Automatic Registration of 2D Histological Sections to 3D microCT Volumes: Trabecular Bone

E. L. Lundin[a,b], M. Stauber[a], P. Papageorgiou[c], M. Ehrbar[c], C. Ghayor[d], F. E. Weber[d], C. Tanner[b], O. Goksel[b]

[a]*Scanco Medical AG, Brüttisellen, Switzerland*
[b]*Computer-assisted Applications in Medicine, ETH Zürich, Switzerland*
[c]*Department of Obstetrics, University Hospital Zurich, Switzerland*
[d]*Oral Biotechnology & Bioengineering, Center of Dentistry, University of Zurich, Switzerland*

Abstract

Histomorphometry and microCT are the two dominant imaging techniques to study bone structure and quality to evaluate repair, regeneration, and disease. These two methods are complementary; where histology provides highly resolved tissue properties on a cellular level in 2D, microCT provides spatial information of bone micro-structure in 3D. For this reason, both of these modalities are commonly used in bone studies. As it is not trivial to combine the images of these two modalities, the two methods are typically applied to different specimens within a study. However, we believe that applying both imaging modalities to the same specimen with a suitable fusion strategy may further strengthen the value of each modality. Therefore, we propose a registration method to align 2D histology slices with a 3D microCT volume, without any prior knowledge of the sectioning direction. In a preprocessing step, bone is extracted from both images. Then, we use a strategy for initializing potential locations, and an iterative approach for searching for an ideal fitting plane using Radon-based rigid transforms and feature-based affine alignments. The algorithm was tested and validated with simulated and real data. For the latter, microCT images of trabecular bone with 76 corresponding histological sections acquired from decalcified and calcified specimens were used. The registration resulted in 94.7% acceptable solutions as defined by a registration orientation error of less than 3°. Average registration accuracy of the acceptable results was 0.6°, leading to a target registration error for our method of 106.3 µm, computed based on landmarks annotated by an observer. This corresponds roughly to 10 pixels in the images; although, the relation to actual visible structures that provide the features to register, is arguably more relevant.

Keywords: Trabecular bone, histology, microCT, 2D-3D image registration

1. Introduction

In basic bone research and in pre-clinical studies, histomorphometry and micro-computed tomography (microCT) are considered to be the two standard techniques to quantify bone microstructure and tissue properties. These two techniques yield complementary information, and thus are often both used to get a detailed insight into the pathogenesis and drug response of metabolic bone diseases [1, 2].

Histomorphometry is a technique, which allows to analyze the bone quality and tissue morphology by the use of few sections from a bone specimen. [3]. It allows for quantification of tissue properties and remodeling mechanisms on a cellular level [2]. However, as the structural information is measured in 2D, models are required to extrapolate the information to 3D [4, 5]. These models may add to experimental uncertainties and may prevent exacting conclusions.

MicroCT is a non-destructive three-dimensional imaging technique to assess the internal structure of an object at very high resolution (e.g. 1-10 µm for a tissue specimen) [6]. It provides excellent contrast of bone to soft tissue, and is typically used to quantify bone mineral density and bone microstructure in three-dimensional space [7]. The strength of this technique is its non-destructive characteristic and its ability to retrieve 3D measures without the need of any underlying model. On the other hand, without any contrast agent, microCT has limited contrast for soft tissues.

Typically, these two techniques are applied on different specimens within a study, where the information from both can only be combined at a statistical level but not in a direct fashion of a one-to-one structural correspondence. However, we believe that fusing these modalities will strengthen the conclusions from their complementary information. An example is the analysis of implant ingrowth in bone (e.g., dental implants or implant screws). In this case, the bone-to-implant contact (BIC) can be very accurately measured in 2D. Although the histological 2D BIC was shown to be in good agreement with the 2D BIC of the corresponding slice extracted from the 3D microCT volume, this 2D measurement does not necessarily represent the 3D BIC. The BIC of a 2D slice is only valid for the given cutting orientation [8], but might be very different orthogonal to this slice. Thus, the *accuracy* of the 2D histological slice perfectly complements the *spatial information* retrieved from 3D microCT. Another example is the determination of biomate-

*Corresponding author
Email address: mstauber@scanco.ch (M. Stauber)

rial ingrowth in drill defect studies. From the histological slices, the remains of biomaterial and newly formed bone can well be differentiated. This information is typically more difficult to retrieve from microCT data. Nevertheless, if the two modalities could be fused, such 3D extrapolation could be performed faithfully, therewith yielding a better picture of the resorption process. A fusion would similarly benefit basic scientific studies, e.g. with rodents, where 2D histological information can be referred within its actual 3D spatial context. For instance, one having extracted a certain cell type (e.g., hematopoietic stem cells) from 2D histology can then infer and test hypotheses on their distance and distribution (e.g., with respect to trabecular network) within registered 3D microCT, which otherwise would not be possible from either modality alone.

In general, finding the position of a histological slice within a 3D volume would strengthen the results derived from the histological section in several ways, e.g., by incorporating the 3D spatial distances or for determining whether the histological section is representative for the remaining tissue. Modality fusion would also help in cases where a particular slice is targeted for sectioning, by validating the cutting position after extracting the section. In the scenarios above and many others, the position of a 2D histological section must be identified within the 3D microCT volume, which is a significant challenge.

The cutting orientation of the histological section with respect to microCT is not easy to determine and often not known for the fusion process. Manually identifying the position and orientation of a histological section in a microCT volume is a non-trivial and tedious task. For this reason, several groups have tried to solve this 2D-3D registration problem in an automatic or semi-automatic way. Museyko *et al.* performed semi-automatic registration of mouse vertebrae and tibiae [9], with standard intensity-based registration techniques available in the Insight Segmentation and Registration Toolkit (ITK) [10]. To overcome the small convergence radius typical for 2D-3D registration, they performed manual initialization of the registration by first aligning three approximately corresponding points in microCT and histology images, which is not trivial and not repeatable. Others used specific specimen information to initialize registration algorithms [11, 12, 13, 8]. Specifically, automatic registration of histology images, sectioned near-perpendicular to the axis of a cylindrically shaped biopsy, was performed for cerebral tissue [11] and jaw bone [12, 13]. This work used the cylinder axis as an initial guess for sample orientation, where feature-based matching was performed with slices perpendicular to this axis. This determines potential point correspondences, from which a plane and corresponding affine transformation could be estimated. Automatic registration has also been performed with intensity-based registration methods for dental implants [8], where the implant axis was calculated from segmented images and thus provided an initial orientation. However, such knowledge of axis from cylindrical biopsies or implants is not available for most other histological studies and thus, a generic and repeatable registration strategy is required for the benefit of many other experimental and clinical configurations.

In this article, we propose a novel method to automatically locate the position of a histological section in the microCT volume, without the need for any initial guess. The proposed method is more general than the previous approaches, as it is not dependent on assumptions concerning sample geometry. Our method uses binarized images as input, which are obtained in a preprocessing step by segmenting stained histology slices. Herein we show examples with H&E and Goldner's Trichrome staining protocols. We conclude that our method enables the combined assessment of histomorphometry and microCT and, as such, it will lead to a better insight into bone metabolic processes.

2. Materials

Nine specimens harvested from porcine vertebrae were used for developing and testing our algorithm. The specimens were scanned with a μCT 100 (Scanco Medical AG, Brüttisellen, Switzerland), before histology sectioning was performed with two different techniques; calcified and decalcified sectioning.

For the microCT measurements, the samples were fixed in 4% (v/v) Formalin for 24 hours and then immersed in phosphate buffered saline and stabilized with polystyrene foam to avoid any motion artifacts. Then, they were scanned at a nominal resolution of 10.7 μm, an energy of 70 kVp, an intensity of 114 μA, and an integration time of 1200 ms.

Decalcified histological sectioning (12 days in 10% EDTA, pH 7.4) was performed on three specimens. Specimens were embedded in paraffine and serial sectioning was performed with section distances ranging from 20 to 200 μm and a section thickness of 4 μm. Ten sections were acquired from each sample and 2 to 8 sections of each sample were further evaluated in this study (remaining sections were considered too distorted from the sectioning process). These sections were stained according to standard H&E staining protocol. Images were acquired with a Zeiss 200M inverted microscope, where pixel resolution was 2.55 μm, as measured with the microscope software (Fig. 1(a)).

Calcified sectioning was performed on six specimens. Specimens were embedded in methylmethacrylate and serial sectioning was performed, with a section distance of 100 μm and section thickness of 10 μm. Ten sections from each sample were prepared and all were used for this study. Staining protocol was the Goldner's Trichrome method [14] (Fig. 1(b)). The resolution of light microscope images was 4-8 μm, as determined from an imaged calibration ruler (not shown).

3. Methods

We propose a registration method that takes segmented images as input and aims to locate the sectioned plane of the histological slice, without any a priori knowledge of the sample orientation for serial sectioning. In general, the sectioning plane is an arbitrary, oblique plane in the microCT volume. The sectioning plane has two rotational and one translational degree of freedom, parametrized in this work in spherical coordinates (r, θ, ϕ). Out-of-plane deformations can be assumed to be negligible [9]. However, sectioning generally causes shearing, thus

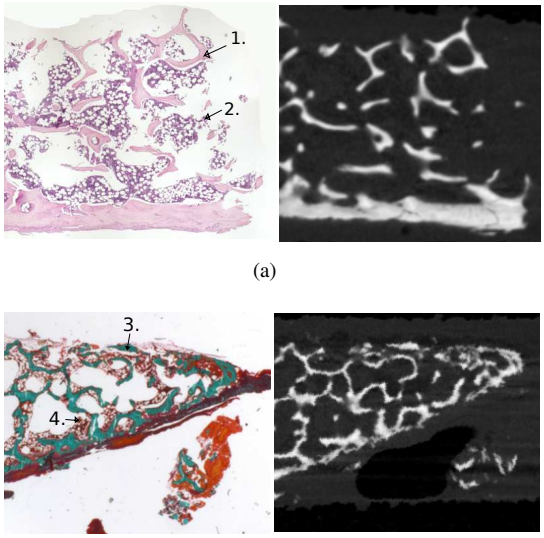

(a)

(b)

Figure 1: (a) H&E stained histology section (left) and corresponding microCT slice (right). 1. Mineralized bone. 2. Hematopoietic cells. (b) Goldner's Trichrome stain (left) and corresponding microCT (right). 3. Mineralized bone. 4. Hematopoietic cells.

the in-plane correspondence between the histology image and the microCT plane has to be also determined to solve the registration problem. We model this in-plane deformation with affine transformation.

The full 2D-3D registration problem with 9 degrees of freedom (DOF) with 3 plane parameters above and 6 for affine shearing is non-trivial, as similarity measures between images would then behave highly non-convex. To boost the registration efficiency and robustness, we first perform an approximate rigid (3 DOF) location alignment between the histology image and microCT planes and then perform a subsequent in-plane affine alignment, in effect decoupling the DOFs. The entire registration is performed in four stages, as sketched in Fig. 2 and summarized next. Thereafter individual methods are described in details.

In the first stage **S1**, several plane normal orientations (given by (θ, ϕ)) are sampled in 3D and candidate plane orientations are determined from a cost function after an intensity-based rigid in-plane registration. In the second stage **S2**, a series of parallel slices are extracted for each of these orientation candidates. For each slice, more expensive affine in-plane registration to the histological image is performed with a feature-based matching method. Feature-based methods match points in images based on some measure of similarity between the neighborhood of the points. Thereafter, a model is fitted to the matching points with a robust regression method, separating points in good agreement with the model, the *inlier points*, from remaining points, the *outliers*. We call the procedure of repeatedly performing feature-based in-plane registration for a series of parallel slices as **Scanning** (Algorithm 1). The resulting 2D inlier points are filtered and their 3D coordinates are stored to form a set of 2D-3D matching points. In the third stage **S3**, a

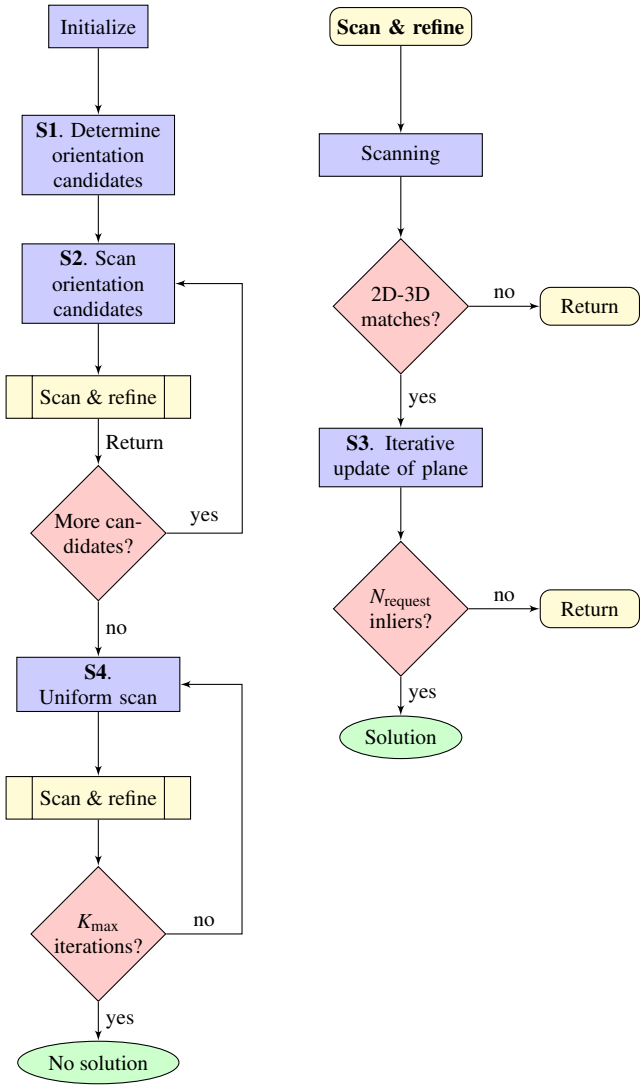

Figure 2: Registration flowchart.

plane is estimated from this set of coordinates. The plane parameters are refined by iterative scanning and plane estimation, until the number of 2D inliers N_i for each estimated plane stops increasing ($N_{i+1} < N_i$). If more than a certain number of inliers (N_{request}) are found, the registration solution is accepted and the algorithm terminates. The fourth stage **S4** is only used if none of the candidate orientations from S1 leads to a registration solution. In this case, random plane orientations are drawn from a uniform distribution and then refined.

3.1. Preprocessing

Before registration, the images are downsampled and segmented as described next.

3.1.1. Adaption of image resolution

To reduce memory requirements and increase registration speed, images are downscaled to two different levels (low/high).

In stage S1, low-resolution images are used, as this is sufficient and efficient for the intensity based similarity metric. Low

resolution can be adapted such that trabeculae are resolved with a few pixels/voxels in width. The optimal *low* resolution thus depends on the average thickness of trabeculae in images. In our study, a fixed resolution of S_{low}=85.6 µm was used (1/8 of the original nominal microCT resolution).

Feature-based matching requires higher resolution for the trabeculae, which we set in our study as S_{high}=21.4 µm (1/2 of the original nominal microCT resolution).

All images were downsampled, after Gaussian smoothing, to these given two physical resolutions.

3.1.2. Image segmentation

We first extract (interpolate) 2D microCT intensity images and afterward segment mineralized bone in microCT images, which can simply be done by employing a global minimum intensity threshold. The threshold was set to 443 mgHA/ccm for high resolution images and to 353 mgHA/ccm for low resolution images due to partial volume effect.

Segmentation of mineralized bone in histological images is a more complex task and depends on the staining protocol. Below, we describe two strategies for the two staining protocols studied.

With the Goldner's Trichrome method, mineralized bone is stained dark green (Fig. 1(b), arrow 3) while osteoid and hematopoieitic cells are stained red (Fig. 1(b), arrow 4). Let I_{r}, I_{g}, and I_{b} denote the RGB intensities of the histological image, and I_{gray} be the gray value intensity converted using the standard luminosity method, i.e. I_{gray}=0.2989I_{r} + 0.5870I_{g} + 0.1140I_{b}. The proposed segmentation strategy is then:

1. Keep (green) pixels with $I_{\text{g}} > I_{\text{r}}$ and $I_{\text{b}} > I_{\text{r}}$
2. Keep pixels below a threshold with $I_{\text{gray}} < T_{\text{gray}}$
3. Remove small connected-components ($< S_{\text{comp}}$ pixels)

where T_{gray} (global threshold) and S_{comp} (size of components) were manually tuned to obtain best trabecular segmentation. S_{comp} was set to a constant value for all images and T_{gray} was constant for images recorded with the same microscope brightness settings. Based on visual inspection, the chosen approach yielded satisfactory segmentation for all histology images.

H&E is a standard stain, where mineralized bone is stained pink (Fig. 1(a), arrow 1) and hematopoietic cells purple (Fig. 1(a), arrow 2). We use the following strategy with an additional red-blue threshold parameter T_{rbRatio} for H&E images:

1. Set T_{rbRatio} = 0.8
2. Keep pixels with I_{r}>I_{g} and I_{b}>I_{g} and $I_{\text{r}}/I_{\text{b}}$>$T_{\text{rbRatio}}$
3. Require $I_{\text{gray}} > T_{\text{gray}}$ and manually tune T_{gray} to the highest value for which all trabeculae are segmented
4. Increase T_{rbRatio} until hematopoietic cells are no longer segmented
5. Remove small connected-components ($< S_{\text{comp}}$ pixels)

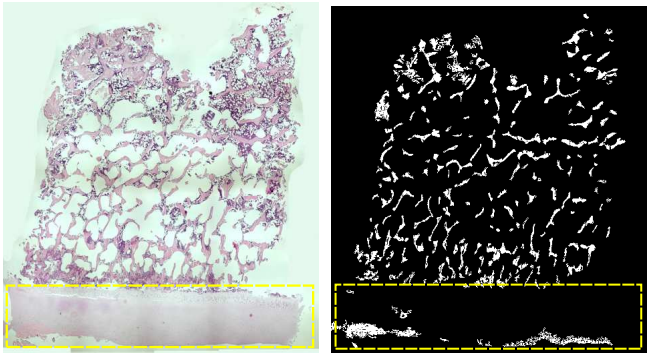

(a) (b)

Figure 3: Segmentation of H&E stained images. (a) Original image (b) Binary image after segmentation. In this specific specimen, there were remains of non-mineralized tissue (indicated by yellow rectangle), which was removed manually.

As brightness conditions varied from image to image for H&E stained sections, thresholds were re-tuned for each image. Since the extracellular matrix, including cartilaginous and non-mineralized bone, also stains pink with the H&E stain protocol, it can be difficult to distinguish these from mineralized bone. Thus, remains of any non-mineralized tissue was removed manually (Fig. 3) on the sections of one sample.

3.2. In-plane registration

Before describing the overall registration algorithm in the next section, we introduce here the two types of in-plane registrations. For both types, it is desired to avoid time-consuming explicit optimization. Therefore, the Radon transform is used for rigid in-plane registration in low resolution. It provides very fast, but approximate, evaluations of the plane parameter search space in stage **S1**, where the plane parameters are initially optimized. It is further used to rotationally pre-align high resolution images before affine matching is performed in the later stages. For affine registration, feature-based methods are used. They efficiently select a distinctive number of points in images that can be used to estimate a model transformation by regression and thus avoid any explicit optimization. They are suitable for high contrast images that contain well-resolved structures, but can only be used for transformation models with a low number of parameters. The problem of affine registration of images containing trabecular bone fulfills these conditions.

3.2.1. Intensity-based rigid in-plane alignment

For a given microCT plane (defined by distance r, angles θ and ϕ), the in-plane rigid alignment to the histological image is determined approximately by alignment of the image centroids $(\bar{x}, \bar{y})$ and determination of the relative rotation from the Radon transform ρ-variances.

Radon transform. The Radon transform R of image I is defined as the line projection of I, for various line orientations

and translations, i.e.

$$R(\rho,\xi)[I(x-\bar{x},y-\bar{y})] = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} I(x-\bar{x},y-\bar{y})\delta(\rho - x\cos(\xi) - y\sin(\xi))\,dx\,dy, \quad (1)$$

where ρ is the perpendicular distance from the line to the origin and ξ is the angle formed between the line and the y-axis. In texture analysis [15], rotation invariance was achieved by using the rotation angle $\beta = \arg\min_\xi d^2\sigma^2(\xi)/(d\xi^2)$, where $\sigma^2(\xi)$ denotes the variance over ρ for angle ξ. Here, we cross-correlate the variances from the two images to determine their relative rotation β_{rel} to each other.

ED metric. The performance of the rigid in-plane alignment is evaluated based on the edge distance (ED) metric [16] between the estimated microCT slice and the binary histological image. Most of the trabeculae of the 2D images will be slightly displaced with respect to the other image after rigid registration. It is therefore important to use a metric like ED, which is sensitive to displaced edges, instead of relying on the conventional segmentation overlap. Edges are extracted with Sobel filtering, forming a set of edge coordinates $\mathcal{E}$. The edge image is distance transformed [17] into a new image I_{dist}. The ED metric between microCT slice with edge coordinates $\mathcal{E}_{\text{CT}}$ and histological image with distance transform $I_{\text{dist,HI}}$ is calculated as

$$c_{\text{CT}\to\text{HI}} = \frac{1}{M}\sum_{i=1}^{M} I_{\text{dist,HI}}(x_i,y_i)^2, \quad (x_i,y_i) \in \mathcal{E}_{\text{CT}}, \quad (2)$$

where M is the number of points in $\mathcal{E}_{\text{CT}}$. In this work, both distances from histology edges to microCT slice edges ($c_{\text{CT}\to\text{HI}}$) and edge distances from slice to histological image ($c_{\text{HI}\to\text{CT}}$) are considered. The final ED metric is the weighted sum $c = (1-\alpha)c_{\text{CT}\to\text{HI}} + \alpha c_{\text{HI}\to\text{CT}}$.

Cost. Since Radon transform cannot resolve flipped correspondences, the ED metric is calculated for four different alignments of the histological image to the microCT slice: for flipped and un-flipped histological image, and for both β_{rel} and $\beta_{\text{rel}} + 180°$ for each flip direction. This results in the four-element cost vector $\mathbf{c}$ and corresponding state vector $\boldsymbol{\beta}$, containing the relative rotations of each state obtained from the Radon transform alignment.

In stages S1 and *Scanning*, the best matching parameters from the four flip alternatives and the corresponding cost f_{cost} are used, where

$$f_{\text{cost}}(r,\theta,\phi) = \min_{i\in\{1..4\}} c_i(r,\theta,\phi)\,. \quad (3)$$

3.2.2. Feature-based affine in-plane alignment

Feature-based in-plane alignment is based on detecting and matching keypoints in high resolution. Inlier points are separated from outliers and thereafter an affine transformation is fitted to the matched inlier keypoints.

Keypoint detection. Keypoints are detected with Harris corner detector [18]. Descriptor vectors are extracted for each keypoint with a simplified version of HoG [19], which skips the non-linear weighting step in gradient orientation binning. This simplification (called bHoG4) is motivated by the fact that gradient orientations in binary images only take on discrete values (with 45° separation). The bHoG4 descriptor is extracted according to the following steps:

1. Compute image derivatives I_{x} and I_{y} by convolving I with the [1 0 -1] and [1 0 -1]^T filter, respectively.
2. Separate positive and negative binary responses p_{x}=(I_{x}>0), p_{y}=(I_{y}>0), n_{x}=(I_{x}<0) and n_{y}=(I_{y}<0).
3. Create blocks of pixels surrounding each keypoint. A block consists of $S_{\text{block}} \times S_{\text{block}}$ cells, that are formed by $S_{\text{cell}} \times S_{\text{cell}}$ pixels. For each cell, the binary gradient responses are summed up over all pixels to create local descriptors with four values.
4. The final descriptor consists of all the local descriptors in a block and is $S_{\text{block}} \times S_{\text{block}} \times 4$ dimensional.

Keypoint matching. Keypoints are nearest-neighbor matched based on their feature descriptors [20]. Descriptor vectors are normalized and compared using their Euclidean distance $d_{j,k}$ (j=1, ..., N_1, k=1, ..., N_2; N_1, N_2 number of keypoints). Nearest neighbor matches are only accepted as matches if $d_{j,k}$<T_{match}, where the parameter T_{match}∈[0, 2]. Additionally, the uniqueness criterion from [21] is applied, which requires d_{j,k_0}<$F_{\text{match}}d_{j,k_1}$ for k_0 being the best match and k_1 the second best match, where F_{match} is a parameter.

Affine transformation fitting. Matched 2D inlier points are initially separated from 2D outlier points with OptimalRANSAC [22] and an affine transformation is estimated. Among RANSAC-like algorithms, OptimalRANSAC was chosen to minimize stochastic effects and optimize robustness.

2D inliers. A slice is considered to contain 2D inliers if the number of inliers after OptimalRANSAC is greater than three and if the scaling of the corresponding affine transformation is limited. The latter is measured by the transformation's singular values $\boldsymbol{\sigma}$ [23] staying within empirically-set lower and upper bounds, i.e. L_i<σ_i<U_i, for both 2D axes i=1, 2. The coordinates of these 2D inliers are returned in set $\mathcal{I}_{\text{2D}}$.

3.3. Overall registration algorithm

We search the microCT slice that results in the greatest number of 2D inlier point correspondences with a binary histological image. The search space of slice extraction is parametrized in spherical coordinates (r,θ,ϕ), where the polar ($\theta \in [0\ 180]°$) and azimuthal ($\phi \in [0\ 360]°$) angles determine the plane normal direction and r the distance from the plane to the origin. However, since it is possible that sections were imaged either from the front or the back (resulting in flipped images), in-plane alignment also considers flips.

3.3.0. Initialization

Preprocessing of histological images is done once at initialization, whereas microCT slices are repeatedly rebinarized as different slices are extracted. The registration is initialized by

computing the Radon transform, distance transform and edges (low resolution), corners and descriptor vectors (high resolution), for both flipped and un-flipped histological image.

3.3.1. S1. Search orientation candidates

In the initial stage of the registration algorithm, candidates for the approximate orientation of sample serial sectioning are determined based on the cost function f_{cost} introduced in Section 3.2.1, i.e. accounting for low-resolution rigid in-plane alignment. Local minima of the cost function are determined with a multi-start approach of many local optimizers. The approximate in-plane alignment of images causes $f_{\text{cost}}(r, \theta, \phi)$ to behave non-smoothly. Thus, the direct search method *patternsearch* [24], which can cope with this, is used as an optimizer.

N_{Opt} repeated optimizations are attempted with different initializations, where each is initialized by drawing N_{Pts} directions $\hat{\mathbf{n}}(\theta, \phi) = [\cos(\phi)\sin(\theta), \sin(\phi)\sin(\theta), \cos(\theta)]$ from a uniform spherical-angular random distribution. The parameter r is drawn from the uniform distribution $r \in [0, R(\theta, \phi)]$, where $R(\theta, \phi) = \hat{\mathbf{n}}(\theta, \phi) \cdot \mathbf{v}_{\text{dim}}/2$ and $\mathbf{v}_{\text{dim}}$ is a vector with the dimensions of the 3D image. The N_{Opt} points with the lowest f_{cost} are chosen as initial points for *patternsearch*.

Finally the N_{Opt} solutions from *patternsearch* are sorted according to increasing f_{cost}. The solutions $f_{\text{cost},i} < F_{\text{start}} f_{\text{cost},0}$ within a margin F_{start} of the lowest cost $f_{\text{cost},0}$ are collected in a set C as the most promising alignment candidates (θ,ϕ,r) and their resulting costs.

Algorithm 1: Scanning

```
function Scan (r, θ, φ, M)
  In/out : M
  Output: R
M_l, R ← 0
c, β ← Rigid2DReg(r, θ, φ)                          // Sec. 3.2.1
f_start ← f_cost(r, θ, φ)
M_l, R ← Add2DInliers(r, θ, φ, c, β, M_l, R)
for s ∈ {−1, 1} do
    r' ← r; k ← 0
    while k < N_stop do
        r' ← r' + sΔ_r
        c, β ← Rigid2DReg(r', θ, φ)                 // Sec. 3.2.1
        f ← f_cost(r', θ, φ)
        M_l, R ← Add2DInliers(r', θ, φ, c, β, M_l, R)
        if f > f_start F_stop then
            k ← k + 1
        end
    end
end
M ← add M_l if size(R) > 2

function Add2DInliers (r, θ, φ, c, β, M_l, R)
  In/out : M_l, R
c, β ← Rigid2DReg(r, θ, φ)                          // Sec. 3.2.1
I_2D ← Affine2DReg(r, θ, φ, c, β)                   // Sec. 3.2.2
M_l ← add I_2D
R ← add r if I_2D non-empty
```

3.3.2. S2. Orientation scanning

For each candidate in C, a series of parallel microCT slices will be extracted: the orientation candidate is scanned in search for points that match between the 2D histology image and the current microCT slice. Slices have plane normals in the direction of the current orientation candidate and are extracted with separation Δr. A scan is started from the r-value corresponding to the local minima of $f_{\text{cost},i}$, for the current candidate i in set C. Slices are then extracted both for increasing and decreasing r-values. For each slice, efficient rigid in-plane registration is performed in low resolution. This provides the rotational in-plane orientation of the slice as well as its cost. Cost values determine when the scan terminates, as outlined in Algorithm 1. After rigid in-plane alignment, affine alignment is attempted, by matching the high resolution slice to histological image with feature-based methods (described in Section 3.2.2). Affine matching is performed for all orientations in $\boldsymbol{\beta}$ corresponding to the ED-costs $\mathbf{c}$, starting with the lowest value, and then progressing to higher values if no 2D inliers were detected. As a scan finishes, the 3D coordinates of all detected 2D inliers will be added to a global set of 2D-3D matches $\mathcal{M}$, provided that 2D inliers were encountered at more than one scan position r for given scan direction. This decreases the likelihood of adding false 2D-3D matches and assures that all points in $\mathcal{M}$ are not contained in the same plane. When a scan results in additional points in $\mathcal{M}$, the registration proceeds to stage S3.

3.3.3. S3. iterative plane estimation

In stage S3, a microCT plane is estimated from the 3D coordinates of the points in $\mathcal{M}$. The plane is iteratively updated by performing repeated scanning and plane estimation, until convergence to plane (r^*, θ^*, ϕ^*), as described in Algorithm 2. (r^*, θ^*, ϕ^*), together with its in-plane affine transform, is accepted as the registration solution provided that its number of inliers $n^* > N_{\text{request}}$, which terminates the registration. Planes are estimated with linear orthogonal regression (Chapter 2 in [23]) together with OptimalRANSAC [22].

3.3.4. S4. uniform scanning

If no registration solutions has been found for any of the orientation candidates from C, the registration enters a stage of uniform scanning of the angular search space (θ, ϕ). Scan directions are drawn randomly from a uniform spherical distribution. These scan directions are pre-scanned with rigid in-plane alignment, in low resolution.

Scanning with the more expensive affine matching will thereafter be performed, as in Algorithm 1, only if slices with $f_{\text{cost}} < F_{\text{start}} f_{\text{cost},0}$ are encountered. Scanning will then start from the r value of the lowest f_{cost}. If 2D-3D inliers are detected, S3 is repeated. New scans are performed until a registration solution is found, or until $K_{\max}$ scan directions were evaluated.

Algorithm 2: S3 Iterative Plane Estimation

function S3 ($\mathcal{M}$)
Output: n^*, (r^*, θ^*, ϕ^*)
$(r_t, \theta_t, \phi_t) \leftarrow$ estimate plane from 3D coord. $\in \mathcal{M}$
$\mathbf{c}, \boldsymbol{\beta} \leftarrow$ Rigid2DReg(r_t, θ_t, ϕ_t) `// Sec. 3.2.1`
$\mathcal{I}_{2D} \leftarrow$ Affine2DReg$(r_t, \theta_t, \phi_t, \mathbf{c}, \boldsymbol{\beta})$ `// Sec. 3.2.2`
$n_t \leftarrow$ number of 2D inliers $(\mathcal{I}_{2D})$
$n^* \leftarrow -1$
while $n^* < n_t$ **do**
  $n^* \leftarrow n_t$
  $(r^*, \theta^*, \phi^*) \leftarrow (r_t, \theta_t, \phi_t)$
  $\mathcal{M}, \mathcal{R} \leftarrow$ Scan$(r^*, \theta^*, \phi^*, \mathcal{M})$
  $(r_t, \theta_t, \phi_t) \leftarrow$ estimate plane from 3D coord. $\in \mathcal{M}$
  $\mathbf{c}, \boldsymbol{\beta} \leftarrow$ Rigid2DReg(r_t, θ_t, ϕ_t) `// Sec. 3.2.1`
  $\mathcal{I}_{2D} \leftarrow$ Affine2DReg$(r_t, \theta_t, \phi_t, \mathbf{c}, \boldsymbol{\beta})$ `// Sec. 3.2.2`
  $n_t \leftarrow$ number of 2D inliers $(\mathcal{I}_{2D})$
end

3.4. Experiments

The registration performance was evaluated for simulated 2D images and for segmented 2D histological images.

3.4.1. Simulated 2D data

2D slices were extracted from the microCT volume to simulate histological images. In total 90 slices were randomly extracted from nine samples. From each sample, three *reference* slices were extracted at the centroids of the microCT images, with plane normal directions along the axes of the principal components of the segmented data. Typical histology slices contain a fair amount of trabecular bone with some structural complexity. To ensure this on simulated microCT slices, we enforced two constraints: that bone appearing on each slice had at least $\geq 25\%$ of the smallest area of the references and had a number of connected components $\geq 50\%$ of the smallest number of the references. First, slices were extracted and their 3D-coordinates stored. Second, 2D in-plane transformations were applied to the extracted slices and to their ground truth 3D-coordinates. Finally the slices were segmented with the same threshold as used for the 3D microCT images.

Three in-plane transformation scenarios were simulated; rigid, affine and local deformations. These three scenarios cover the most important deformations occurring to the 2D histological slice during the sectioning process, to a varying degree of registration difficulty. For the rigid scenario, random rotations were used for simulating transformations. For the affine scenario, images were anisotropically scaled along a randomly selected coordinate frame. The singular values s_1 and s_2 of the transformation, defining the anisotropic scale changes, were sampled from uniform distributions $s_1 \in [0.93, 1.12]$ and $s_2 \in [0.75, 1.08]$. These intervals correspond to the range of singular values of the affine transform of gold standard microCT planes to histological images, with an extra margin of 0.05.

For the local deformation scenario, slices were simulated by picking five random points within the convex hull of the segmentation and assigning them random displacements $\mathbf{d}=[d_x, d_y]$, with d_x, d_y from $\mathcal{N}(0, \sigma_d^2)$. σ_d was selected randomly for each slice, from the uniform distribution $\sigma_d \in [0, 85]$ µm. A dense displacement field was calculated from these point displacements using thin plate splines [25] and inverse distance interpolation [26] was performed.

The simulated slices were then registered to the 3D microCT images and the results compared to their known ground truth 3D coordinates.

Table 1: Registration parameters, which are either in pixels (marked by $^{\#}$) or unitless.

Parameter	Value	Function
α	0.7	ED weight
$S_{\text{cell}}{}^{\#}$	7	bHoG4
$S_{\text{block}}{}^{\#}$	6	bHoG4
T_{match}	0.4	Nearest neighbor matching
F_{match}	0.6	Uniqueness criterion
$T_{\text{tol}}{}^{\#}$	10	Model tolerance
$T_{\text{acc}}{}^{\#}$	8	Pruning tolerance
$[L_1, U_1]$	[0.8, 1.5]	2D singular value
$[L_2, U_2]$	[0.6, 1.2]	2D singular value
N_{Pts}	60	S1, initial points
N_{Opt}	20	S1, local optimizers
F_{start}	3	S2, S4, scanning
F_{stop}	3	S2, S4, scanning
$\Delta_{\text{r}}{}^{\#}$	2	S2, S4, scan interval
N_{stop}	4	S2, S4, scan termination
N_{request}	10	S3, 2D inliers
K_{max}	500	S4

3.4.2. Real 2D data

For the real data (acquired as described in Section 2), performance was assessed with respect to semi-automatically determined gold standard microCT slices. For this, an approximate direction of the histology serial sectioning was identified manually. The manual approach was chosen to have a completely independent approach from our method. This direction was then scanned for matching features using a different feature-matching method called SURF [21], to avoid bias in the determination of the gold standard planes. For registration of 2D images, SURF is a powerful method, as it can handle both scale and rotational changes. However, as the transformation space is less constricted, it also results in a higher number of false matches, which may degrade performance in 2D-3D scenarios. The gold standard plane was found by plane-estimation from the detected inlier points. Results were confirmed visually. For each section, five point correspondences between the histological image and microCT gold standard plane were determined manually.

3.5. Quantification of registration errors

Registration results have been evaluated based on target registration error (TRE) [27], dense registration error (DRE), and the angular error Ψ^* from the registered plane normal to the

gold standard or ground truth plane normal. TRE and DRE are both defined between images as the mean Euclidean distance of registered landmarks to corresponding gold-standard locations. For TRE, these are relatively few locations at selected points of interest, while for DRE the locations are regularly sampled within a region of interest. In this study, TRE or DRE error distributions are summarized by their median, mean, and 95%tile values over the images. TRE and DRE assess general registration accuracy, both in-plane and out-of-plane, whereas Ψ^* only assesses the out-of-plane error.

For the simulated data, DRE was evaluated for all pixels within the convex hull of the 2D segmentation. For the real data, TRE was calculated for the five manually selected point correspondences per section.

The measure for initial deformation of the simulated in-plane transformations (DRE_I) and the non-rigid registration error at the gold standard plane (TRE_{GS}) were also determined. DRE_I is based on the DRE between the transformed (affine or locally deformed) and the rigid slice coordinates. TRE_{GS} was measured as the error of the manual landmark annotations after rigid registration of the histological image to the gold standard microCT plane. Thus, TRE_{GS} represents the combined amount of affine and local deformations inflicted onto the section during the histology process. The component due to local deformations (denoted as $TRE_{GS,deform}$) was determined by the TRE after affinely registering the histological image to the gold standard microCT plane.

A registration was classified as not finished if it did not reach a registration solution within K_{max} iterations in stage S4. All registrations with $\Psi^*<3°$ were considered acceptable.

3.6. Registration parameters

All registration parameters are listed in Table 1. The additional parameters used for OptimalRANSAC were taken directly from [22]. The parameters S_{cell} and S_{block} for key-point description, T_{tol} and T_{acc} as model and pruning threshold for robust affine model estimation and the scan interval Δ_r, are all given in pixel units. They all depend on the mean trabecular thickness (Tb.Th). In our study, they were determined empirically from a few images. For future use of this algorithm, they need to be scaled to image resolution. For the scan factor parameters, we choose $F_{start} = F_{stop} = 3$ based on a few trial registrations before the full registration experiments were performed. These factors relate rigid in-plane registration cost of the deepest minima found in **S1** to the expected cost of the true solution. Algorithm robustness as well as registration times increase with increased values of F_{start} and F_{stop}. To the remaining parameters, the algorithm is less sensitive.

4. Results

The registration algorithm has been implemented in C++ on OpenVMS and tested on a HP rx2800 i2 machine.

Four example registrations and the landmarks for TRE calculation are shown in Figure 4. The quantitative results for the simulations and the real histology data are summarized in Table 2. Only one simulated and two real cases did not finish. 95% of all images led to "acceptable" results ($\Psi^*<3°$). Histology sections were registered with an accuracy of TRE = 106 μm and $\Psi^* = 0.6°$ on average. The plane normal was generally more accurately registered for the simulations (mean $\Psi^*\leq 0.3°$).

For the simulated data, rigidly transformed 2D images registered with sub-voxel accuracy (mean DRE = 13.5 μm, S_{high} = 21.4 μm). In Figure 5, the relationship between the initial non-rigid displacement (DRE_I) and the registration accuracy is shown for the affinely and locally deformed simulated data. The applied initial deformation was much larger for the affinely distorted images (mean DRE_I = 240 μm) than for the local deformations (mean DRE_I = 81 μm). Results show that the accuracy of the algorithm is rather independent of the degree of initial affine displacement, whereas larger local deformations result in larger registration errors (Figure 5(a)). The plane normal could be determined reasonably well regardless of the amount of initial deformation (Fig. 5(b)).

In Figure 6, distributions of the registration errors (6(a),6(c)), relative registration errors (6(b)) and registration times (6(d)) for 16 H&E and 60 Trichrome histological images are shown. The relative error is measured as the ratio of TRE to the initial error due to local deformations of the histology slide ($TRE_{GS,deform}$). On average, this ratio is close to one (mean TRE = 106 μm vs. mean $TRE_{GS,deform}$ = 99.8 μm), indicating that the local deformations of the histological slices, which were not compensated, may greatly contribute to TRE.

5. Discussion

In this paper, we present a new method to automatically register 2D histological images to 3D microCT volumes, for highly structured objects such as trabecular bone. With our method, the plane of histological sectioning can be found automatically, without any restrictions on the 3D plane parameter space. Previous automatic methods only compensated for small variations around a pre-determined search space, namely 1D translations along a cylindrical axis [12, 11, 13] or 1D rotations around an implant axis [8]. Thus the proposed method is more general than previous methods, as it makes no assumptions on sample geometry and requires no manual initial localization.

This new method registers slices with sub-voxel accuracy for rigidly transformed images, demonstrating an optimal performance of the algorithm for non-distorted 2D images. To test the sensitivity of the algorithm with respect to image distorsion, test cases with affine and locally distorted 2D images were prepared. These two simulation models were chosen to address alterations, which typically occur during the preparation process of histological slices. The average error Ψ^* with respect to the plane normal was found to be 0.2°–0.3°, which is in the same range as for non-distorted slices (0.2°). Thus, the plane normal can be identified with high accuracy, independent of image distortions. Furthermore, in-plane registration reached an accuracy below 3 voxels (< 64 μm). This is a very high accuracy, although the DRE is four times larger as for the rigid simulations. This is probably caused by errors in the localization of

Table 2: Registration results: total (T) number of 2D images, finished (F) and acceptable (A, Ψ^*<3°) registrations in percentage of the total number. Accuracy results for acceptable registrations are summarized by median, mean and 95%tile. Dense registration errors (DRE) for simulated data are compared to average non-rigid initial displacements DRE_I. TRE for real data is compared to TRE after rigid transformation to the gold standard slice (TRE_{GS}). Run times of finished registrations are presented as percentage of all registrations within stated time intervals.

	T	F [%]	A [%]	Ψ^* [°]			DRE, TRE [µm]			DRE_I, TRE_{GS} [µm]			Time [%]		
				med	mean	95%	med	mean	95%	med	mean	95%	<0.5h	0.5-2h	≥2h
SimRigid	90	100.0	97.8	0.1	0.2	0.6	12.1	13.5	25.4	0.0	0.0	0.0	86.7	8.9	4.4
SimAffine	90	98.9	94.4	0.2	0.3	1.0	62.0	56.4	73.9	225.9	240.1	504.2	82.2	11.1	5.6
SimDeform	90	100.0	96.7	0.1	0.2	0.6	36.8	53.9	168.1	63.1	80.9	253.4	85.6	12.2	2.2
RealH&E	16	93.8	93.8	0.4	0.5	1.3	85.1	96.8	210.8	141.5	166.4	298.1	93.8	0.0	0.0
RealTrichrome	60	98.3	95.0	0.5	0.6	1.8	105.1	108.8	198.6	342.6	370.2	704.7	73.3	15.0	10.0
RealTotal	76	97.4	94.7	0.5	0.6	1.6	98.6	106.3	204.9	317.6	327.3	674.6	77.6	11.8	7.9

matched keypoints in non-rigidly distorted images, resulting in a less accurate transformation.

The algorithm reliably identifies the corresponding microCT slice for both H&E and Goldner's Trichrome stained histological images. For 95% of all slices, our registration method found an acceptable solution with a mean plane normal error $\Psi^* \leq 0.6°$ and a mean target registration error TRE = 106.3µm. This is a very high accuracy in comparison to previous methods [12, 13], where registration errors for cylindrical biopsies were quantified based on the distance between the registered slice and a manually determined slice. Chicherova et al. reported 75% good matches, with a mean registration error of 250 µm [12]. Stalder et al. stated that more than 75% of their sections registered with an accuracy of <100 µm [13]. Their measurement is a 1D error between planes only, while our mean error of 106.3 µm also includes in-plane errors caused by non-affine deformations

The errors in real histology slices were increased as compared to the simulation errors, due to several complicating factors. First, having acquired by a different modality, imaging noise and (trabecular) segmentation inaccuracies affect the registration process. Second, there can be alterations in histological images during sectioning such as structure folding, loss of structures as well as affine and non-affine deformations. Even if our algorithm can identify the correct plane in the microCT data and compensate for affine deformations, these other artifacts are reflected in the TRE values. Third, there are uncertainties in determining the gold standard slices. As shown in Figure 6(b), there were registrations with smaller TRE than the initial local deformations of the gold standard slice ($TRE_{GS,deform}$), indicating that some registrations were potentially even better than the gold standard affine estimation.

The strength of the proposed algorithm is its ability to register a 2D histology slice to the corresponding 3D microCT volume without a prior position estimation to a high accuracy and in reasonable time. For 77.6% of acceptable solutions, the solution was found in under 30 minutes. This is in a similar range as computation times of other automated solutions, that were reported to be in the range from 4 [12] to about 45 minutes [13]. However, computation time depends on many factors such as size of input data, hardware and level of parallelization, and thus it is difficult to perform a direct comparison. In any case, our approach is significantly faster than manual registration, which was estimated to take up to eight hours [12, 11, 13].

Despite all promising results, this algorithm has some limitations and could further be improved. TRE was higher for real images (106.3µm) than simulated (56.4 µm for affine simulations). For increased accuracy, the algorithm could be extended such that a registration, which allows local deformations, is started from the registration solution of the current algorithm.

About 5% of real data registered poorly. One reason for this is that the algorithm did not always reach a registration solution. For real data this happened twice and was due to erroneous estimation of the relative rotation β_{rel} with the Radon transform ρ-variance. Consequently no 2D inliers could be found even though the algorithm had estimated the correct plane relatively early in the process. A rotation-invariant feature descriptor could compensate for errors in the rotation estimation, but would also likely result in an increased number of false 2D-3D matches, since the possible in-plane transformations are less confined. The second cause of these poor registration results is that the algorithm sometimes converged to solutions that were only locally optimal. This is a standard concern for complex optimization problems and our algorithm still showed a higher percentage of acceptable results than previous methods [12, 13].

Finally, a partial limitation of the algorithm is that it requires binary images as input. In this study, we proposed an approach to segment histology slices stained with H&E and Goldner's Trichrome protocols. With this approach, images could be manually segmented in only a few minutes per image. This process could further be optimized and automated to fully automate the whole workflow. However, optimizing the segmentation procedure was beyond the scope of this work. Additionally, this limitation could partially be seen as an advantage, as all histology slices that can be segmented can also be registered.

In summary, we presented a new method for accurate and fast registration of histological slices to microCT images. The method does not require any a priori knowledge of specimen geometry. It expects binary images as input and thus, will work for any stained histology image, that can reasonably be segmented. This reduces manual work to segmenting images, a task, that could be automated as well. Thus, as compared to manual registration, which may take hours of tedious work, our

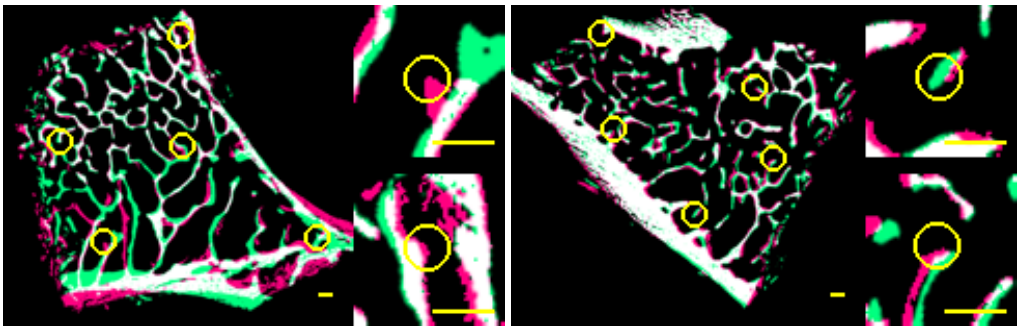

(a) (b)

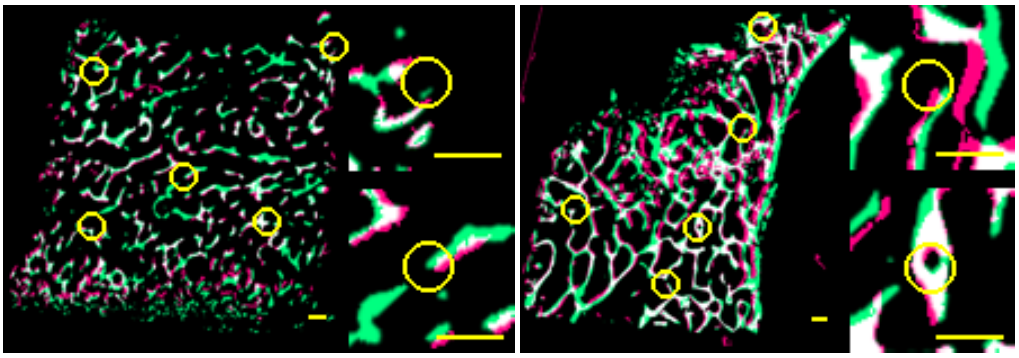

(c) (d)

Figure 4: Registration examples. (a): mean TRE = 220 µm, $\Psi^* = 0.23°$. (b): mean TRE = 123 µm, $\Psi^* = 1.03°$. (c): mean TRE = 122 µm, $\Psi^* = 0.26°$. (d): mean TRE = 74 µm, $\Psi^* = 0.14°$. Pink: histological image, green: microCT slice, white: overlap. Centers of yellow circles mark locations on histological images used for TRE calculation. Horizontal yellow scale bars (bottom right) measure 500 µm. Inserts: (upper) point with largest registration error, (lower) point with median registration error.

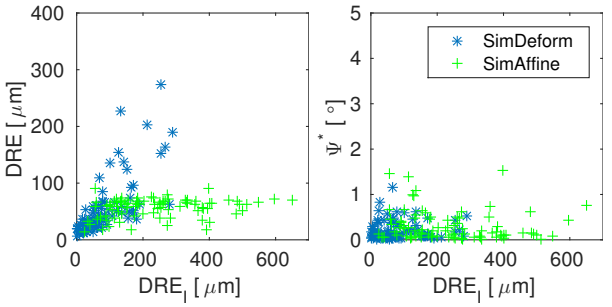

(a) (b)

Figure 5: Acceptable registration results for simulated, affine and locally deformed, 2D images as a function of the introduced mean initial non-rigid 2D displacement DRE_I. (a) DRE vs. DRE_I. (b) Ψ^* vs. DRE_I.

method identifies corresponding slices mostly in less than half an hour fully automatically. We believe, that this enables the combined evaluation of histology and microCT on a larger scale and that fusing this information will help to gain new insight in the bone micro-structural and metabolic changes.

Acknowledgment

We would like to thank Alexandre Tchouboukov and Ana Perez for preparing the calcified histology sections and Esther Kleiner for preparation of the decalcified sections. Funding was partially provided by the Swiss National Science Foundation (SNSF).

[1] D. Chappard, N. Retailleau-Gaborit, E. Legrand, M. Baslé, M. Audran, Comparison insight bone measurements by histomorphometry and µct, American Society for Bone and Mineral Research 20 (7).

[2] C. Moreira Kulak, D. Dempster, Bone histomorphometry: a concise review for endocrinologists and clinicians, Arq Bras Endocrinol Metab.

[3] D. W. Dempster, J. E. Compston, M. K. Drezner, F. H. Glorieux, J. a. Kanis, H. Malluche, P. J. Meunier, S. M. Ott, R. R. Recker, a. M. Parfitt, Standardized nomenclature, symbols, and units for bone histomorphometry: A 2012 update of the report of the ASBMR Histomorphometry Nomenclature Committee, Journal of Bone and Mineral Research 28 (1) (2013) 2–17. doi:10.1002/jbmr.1805.

[4] A. Parfitt, C. Mathews, A. Villanueva, M. Kleerekoper, B. Frame, D. Rao, Relationships between surface, volume, and thickness of iliac trabecular bone in aging and in osteoporosis. implications for the microanatomic and cellular mechanisms of bone loss, J Clin Invest 72 (4).

[5] W. Whitehouse, The quantitative morphology of anisotropic trabecular bone, J Microsc 101.

[6] M. Stauber, R. Müller, Micro-computed tomography: A method for the non-destructive evaluation of the three-dimensional structure of biological specimens, in: Methods in molecular biology (Clifton, N.J.), Vol. 455, 2008, pp. 273–292. doi:10.1007/978-1-59745-104-8_19.

[7] M. L. Bouxsein, S. K. Boyd, B. a. Christiansen, R. E. Guldberg, K. J. Jepsen, R. Müller, Guidelines for assessment of bone microstructure in

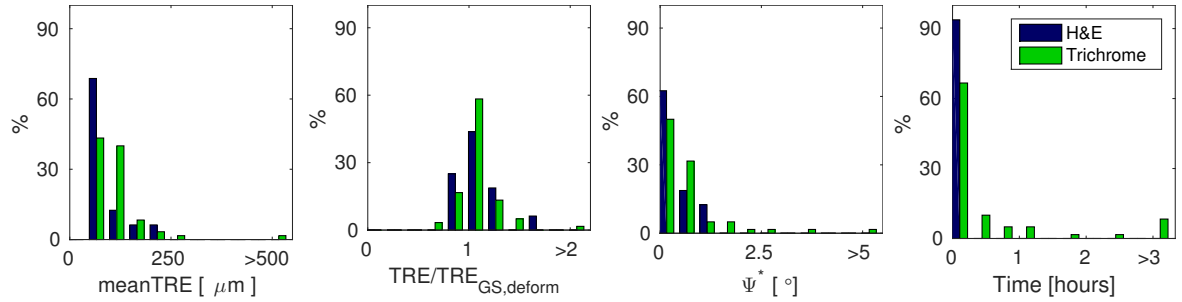

Figure 6: Distribution of registration results of histological images to microCT. (a) Mean TRE per histological image. (b) Ratio of mean TRE to initial mean displacement due to local deformations ($TRE_{GS,deform}$). (c) Angle between registration result and gold standard plane. (d) Registration time.

rodents using micro-computed tomography, Journal of Bone and Mineral Research 25 (7) (2010) 1468–1486. doi:10.1002/jbmr.141.

[8] K. Becker, M. Stauber, F. Schwarz, T. Beißbarth, Automated 3D-2D registration of X-ray microcomputed tomography with histological sections for dental implants in bone, using chamfer matching and simulated annealing, Computerized Medical Imaging and Graphics 44 (2015) 62–68.

[9] O. Museyko, R. M. Marshall, J. Lu, A. Hess, G. Schett, M. Amling, W. A. Kalender, K. Engelke, Registration of 2D histological sections with 3D micro-CT datasets from small animal vertebrae and tibiae, Computer Methods in Biomechanics and Biomedical Engineering 18 (15) (2015) 1658–1673.

[10] L. Ibanez, W. Schroeder, L. Ng, J. Cates, The ITK software guide, Kitware, Inc., New York, 2005.

[11] A. Khimchenko, H. Deyhle, G. Schulz, G. Schweighauser, J. Hench, N. Chicherova, C. Bikis, S. E. Hieber, B. Müller, Extending two-dimensional histology into the third dimension through conventional micro computed tomography, NeuroImage.

[12] N. Chicherova, K. Fundana, B. Müller, P. C. Cattin, Histology to μCT data matching using landmarks and a density biased RANSAC, in: Medical image computing and computer-assisted intervention (MICCAI), Vol. 17, 2014, pp. 243–250. doi:10.1007/978-3-319-10404-1_31.

[13] A. K. Stalder, B. Ilgenstein, N. Chicherova, H. Deyhle, F. Beckmann, B. Müller, S. E. Hieber, Combined use of micro computed tomography and histology to evaluate the regenerative capacity of bone grafting materials, International Journal of Materials Research 105 (7) (2014) 679–691. doi:10.3139/146.111050.

[14] M. Helfrich, S. Ralston (Eds.), Bone Research Protocols, Humana Press, Totowa, New Jersey, 2003, Ch. 24. Studies of Local Bone Remodeling, pp. 345–351.

[15] K. Jafari-Khouzani, H. Soltanian-Zadeh, Radon transform orientation estimation for rotation invariant texture analysis, IEEE Trans Pattern Anal. Mach. Intell. 27 (6) (2005) 1004–1008.

[16] G. Borgefors, Hierarchal chamfer matching: A parametric edge matching algorithm, IEEE Transactions of Pattern Analysis and Machine Intelligence 10 (6).

[17] P. Felzenswalb, D. Huttenlocher, Distance transform of sampled functions, Theory of Computing 8 (2012) 415–428.

[18] C. Harris, M. Stephens, A combined corner and edge detector, in: Proceedings of the 4th Alvey Vision Conference, 1988, pp. 147–151.

[19] N. Dalal, B. Triggs, Histogram of oriented gradients for human detection, in: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Vol. 1, 2005, pp. 886–893.

[20] D. Lowe, Distinctive image features from scale-invariant keypoints, IJCV 60 (2004) 91–110.

[21] H. Bay, T. Tuytelaars, L. Van Gool, SURF: Speeded up robust features, in: A. Leonardis, H. Bischof, A. Pinz (Eds.), European Conference on Computer Vision, Springer-Verlag Berlin Heidelberg, 2006, pp. 404–417.

[22] A. Hast, J. Nysjö, A. Marchetti, Optimal RANSAC - towards a repeatable algorithm for finding the optimal set, Journal of WSCG 1 (2013) 21–30.

[23] A. Björck, Numerical Methods in Matrix Computations, Vol. 59 of Texts in Applied Mathematics, Springer, 2015, Ch. 1, p. 32.

[24] R. Hooke, T. Jeeves, Direct search solution of numerical and statistical problems, J. ACM.

[25] F. L. Bookstein, Principle warps: Thin-plate splines and the decomposition of deformations, IEEE Transactions on Pattern Analysis and Machine Intelligence 11 (1989) 567–585.

[26] D. Shephard, A two-dimensional interpolation function for irregularly-spaced data, in: Proceedings of the 1968 ACM National Conference, 1968, p. 517–524.

[27] C. R. Maurer, J. M. Fitzpatrick, R. L. Galloway, M. Y. Wang, R. J. Maciunas, G. S. Allen, The accuracy of image-guided neurosurgery using implantable fiducial markers, in: Computer Assisted Radiology, Berlin: Springer-Verlag, 1995, pp. 1197–1202.