HAL
open science

Capturing the Multiscale Anatomical Shape Variability with Polyaffine Transformation Trees

Christof Seiler, Xavier Pennec, Mauricio Reyes

To cite this version:

Christof Seiler, Xavier Pennec, Mauricio Reyes. Capturing the Multiscale Anatomical Shape Variability with Polyaffine Transformation Trees. Medical Image Analysis, 2012, 16 (7), pp.1371-1384. 10.1016/j.media.2012.05.011 . hal-00813866

HAL Id: hal-00813866

https://inria.hal.science/hal-00813866v1

Submitted on 26 Apr 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Capturing the Multiscale Anatomical Shape Variability with Polyaffine Transformation Trees

Christof Seiler[a,b], Xavier Pennec[a], Mauricio Reyes[b]

[a] *Institute for Surgical Technology and Biomechanics, University of Bern, Switzerland*
[b] *Asclepios Research Group, INRIA Sophia Antipolis, France*

Abstract

Mandible fractures are classified depending on their location. In clinical practice, locations are grouped into regions at different scales according to anatomical, functional and aesthetic considerations. Implant design aims at defining the optimal implant for each patient. Emerging population-based techniques analyze the anatomical variability across a population and perform statistical analysis to identify an optimal set of implants. Current efforts are focused on finding clusters of patients with similar characteristics and designing one implant for each cluster. Ideally, the description of anatomical variability is directly connected to the clinical regions. This connection is what we present here, by introducing a new registration method that builds upon a tree of locally affine transformations that describes variability at different scales. We assess the accuracy of our method on 146 CT images of femurs. Two medical experts provide the ground truth by manually measuring six landmarks. We illustrate the clinical importance of our method by clustering 43 CT images of mandibles for implant design. The presented method does not require any application-specific input, which makes it attractive for the analysis of other multiscale anatomical structures. At the core of our new method lays the introduction of a new basis for stationary velocity fields. This basis has very close links to anatomical substructures. In the future, this method has the potential to discover the hidden and possibly sparse structure of the anatomy.

Keywords:
parametrization of diffeomorphisms, shape statistics, multiscale,

Email address: `christof.seiler@istb.unibe.ch` (Christof Seiler)

hierarchical, trees, Log-Euclidean polyaffine transformations, locally affine registration, polyaffine registration, Log-Demons registration

---

1. Implant Design for Mandible Fractures

1.1. Clinical Problem

Mandibular fractures most commonly result from facial trauma, with close to half of the patients requiring surgical repair Ellis et al. (1985). A majority of 75% of fractures occur in males aged between 20 and 30 Ellis et al. (1985); Moore et al. (1985), and are often caused by physical assault. Other causes of fractures include motor vehicle collisions, falls, and sports-related injury Craig et al. (2008); Kovan (2008). For all these cases surgical repair proves most effective, with the goal of recovering the anatomical structure prior to the injury and thus restoring normal function. To reach this goal the surgeon places wires or implants at the fracture site, so that the natural fusion of separated bone pieces restores the prior structure Fedok et al. (1998) as closely as possible. A correct repair aligns teeth for food intake, and restores the patient's aesthetics.

In Urken et al. (1991), the authors propose a classification scheme for mandibles based on regions according to anatomical, functional and aesthetic considerations. The online register www.aofoundation.org uses the same classification (Fig. 1) to guide surgeons through the major steps of mandible reconstructive surgery ranging from diagnosis, selection of the optimal surgical approach, to aftercare treatment. According to Moore et al. (1985), the rate of fracture incidence for each of the classified region is as follows: "Symphyseal and parasymphyseal" region 19%, "body" 24%, "angle and ramus" 40%, "condylar process and head" 16%, and "coronoid" 1%. In addition to this classification scheme, the mandible can be subdivided even further into smaller regions, e.g. one region for each tooth. In the image space, this subdivision can theoretically be performed up to the voxel level, where coarser levels enclosed finer ones, representing a hierarchy of regions that can be organized in a tree-like fashion. As the regions become more fine, it is harder to find a consensus among clinicians on the size, shape and location of the region.

Recent work on biomechanical analysis of implants indicates that geometry and topology of implants are crucial to fracture stability. In Lovald et al. (2009), it was shown that implants optimized for the "body" region, (Fig.

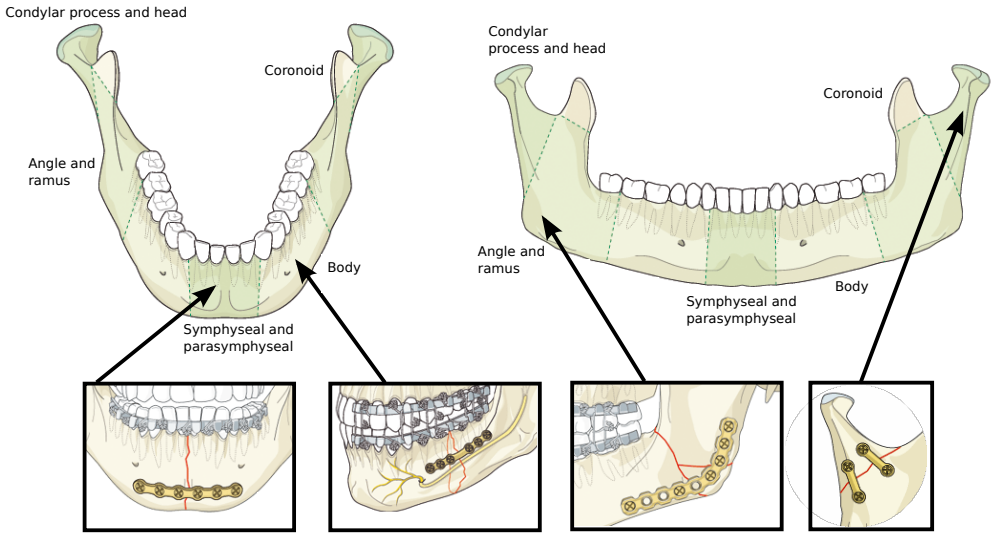

Figure 1: Subdivision of mandible into anatomical regions proposed by the AO foundation to classify fractures for reconstructive surgery. Implants at four different anatomical sites are shown. Images source: www.aofoundation.org.

1), have fracture strain of 69% to 59% and implant stress of 34% to 27% with respect to smaller standard implants, while minimizing patient intrusion by saving 55% of implanted volume of larger standard implants. The same authors presented results on implant optimization of the "symphyseal and parasymphyseal" region with similar results Lovald et al. (2010). In Cervantes et al. (2012), flexible implants that allow the surgeon to adjust the geometry after bone fixation are presented. Instead of pre-manufacturing patient specific implants, there have been several works on population-based designs, for femurs Kozic et al. (2010); Bou-Sleiman et al. (2011); Bonaretti et al. (2011) and mandibles Metzger et al. (2011); Bou-Sleiman et al. (2012). In these approaches, the population is stratified into several sub-populations according to morphological differences or meta-information (e.g. gender, age, etc.). None of the aforementioned population-based approaches consider the bone mineral density in addition to the surface geometry, even though Lovald et al. (2009, 2010) reported its importance to achieve optimal design and placement of screws.

1.2. Methodological Framework

The economical cost of patient specific implants makes this approach impracticable at the moment, hence we focus on an intermediate goal: population-based implant design. Common key steps to population-based design Kozic et al. (2010); Metzger et al. (2011); Bou-Sleiman et al. (2011, 2012) are registration to capture shape variability as encountered in a population and statistical analysis of the registration results, performed subsequently and independently. As mentioned, registration approaches for implant design should consider not only bone surface information, but also volumetric information describing the bone density distribution, which is needed to compute best location and orientation of the placement of screws to fixate implants.

It is a common practice to use principle component analysis (PCA) for dimensionality reduction and extraction of main modes of variation from deformations obtained through non-linear registration. However, due to the global nature of PCA only ad-hoc heuristics to decompose the anatomical shape into localized regions (as shown in Fig. 1) of interest are available, further it is unclear how to interpret the linear combination of different mixtures of shape effects. We therefore conclude that an intelligible link between PCA shape models and implant design for specific regions is missing.

M-reps Pizer et al. (2003); Siddiqi and Pizer (2008) offer remedy for these kind of issues, by modeling surface shapes locally, thus providing a better link between model and clinically motivated regions. Focus on these regions of interest, as opposed to a full non-linear registration, potentially increases the robustness (small number of parameters) and intelligibility (direct connection between regions in a clinical context and shape modeling) of the registration. Unfortunately, due to the fact that the representation is surface-only, volumetric information cannot be included into the analysis to reconstruct bone density distributions.

Therefore we focus on another approach that allows for region-based local description of shape and volumetric information: locally affine transformations, also called polyaffine transformations. Polyaffine transformations were introduced in Arsigny et al. (2005) to fuse locally rigid and affine transformations into a diffeomorphism, and revisited in Arsigny et al. (2009) to obtain faster numerical algorithms. An efficient registration algorithm using approximations of polyaffine transformations was presented in Commowick et al. (2008). In Martìn-Fernández et al. (2009) an extension to articulated structures was presented, which considers weights (defining the influence of each region) fixed at landmark positions along a manually defined skeleton.

Other recent work includes Zhuang et al. (2010), where a Locally Affine Registration Method (LARM) is developed for cardiac MR images. LARMS works by fusing the affine transformations directly, which can cause non-diffeomorphic transformations. These non-diffeomorphic cases are prevented through two additional regularization steps. The authors define regions that are important substructures of the heart and use LARM as an initialization to robustify the subsequent non-linear registration step. In Seiler et al. (2011b), we presented a polyaffine-regularized log-demons algorithm for femur bone registration with manually fixed weights. However, we believe that defining the regions for the locally affine deformations should not be left to the user, because this is subjective and reduces reusability. This becomes even more evident in the case of a multiscale representation of the geometry, where the definition of regions and their division process is not straightforward and time consuming.

To consider more complex shapes and foster reusability, Buerger et al. (2011) presented a multiscale approach with affine regions defined using a data-driven approach. The method splits rectangular shaped regions, which are aligned along the image directions, only if certain conditions are met. Taquet et al. (2011) iteratively optimizes between affine parameters and anchor positions (center of regions) estimation, through an expectation maximization approach. The weights are estimated with a Kriging estimator instead of the usual Gaussian functions. In Zhang et al. (2011), Log-Euclidean Polyaffine Transformations (LEPT) are employed to register multi-modal cardiac sequences in an elaborate scheme. The algorithm iteratively adds new uniform Gaussian weighted regions with different spatial position (mean parameter) and influence (variance parameter) until a mutual information-based criterion is satisfied. The authors in Freiman et al. (2011) showed how the standard diffeomorphic demons (not the log-demons) could be used to enforce an inhomogeneous regularization using a local affine fitting at each voxel of the image. Even though these methods are very promising, the link between the clinical regions (Fig. 1) and the regions found by these algorithms, is either constrained by aligned rectangular shaped regions or spherical weights Buerger et al. (2011); Zhang et al. (2011), or produce an intractable number (around 500 or one region per voxel) of regions Taquet et al. (2011); Freiman et al. (2011).

We present here a new approach with emphasis on interpretability of regions in terms of the clinical requirements (as defined in Fig. 1), by introducing a hierarchical multiscale tree structure that is motivated by the nature

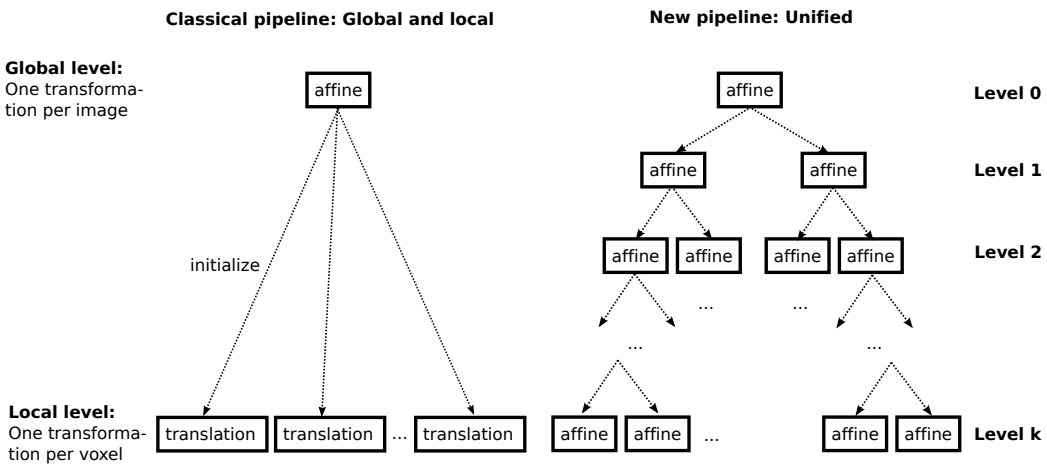

Figure 2: Comparison of classical and newly proposed image registration pipeline. In most classical approaches, the non-linear registration method is initialized with an affine registration performed in advance, clearly separating the two steps. In our new method we propose to move from global to local transformations by subdividing the images into smaller regions, thus providing a smooth transition, while preserving diffeomorphic deformations.

of the mandible anatomy, where regions are ordered and interact with each other in a way tractable for human understanding. Further, since we formulate a data-driven approach, our presented algorithm can be applied to other anatomical structures without modification. To accomplish this we contribute the following points:

1. Linear projection of Stationary Vector Fields (SVF) onto the space of polyaffine transformations (Section 2).
2. Tree structured polyaffine transformations (Section 3).
3. Anatomy-driven definition of regions using a hierarchy of oriented bounding boxes (Section 4).
4. Efficient estimation of transformation parameters with the log-demons algorithm (Section 5).
5. Unified algorithmic registration framework for a "continuous" subdivision of deformations across scales, see Fig. 2 for an illustration.

A preliminary version of this work was presented at SPIE 2011 Seiler et al. (2011b) and MICCAI 2011 Seiler et al. (2011a) conferences. Additionally, in

this paper, we extend the method from surface-based to image-based definitions of regions and provide a statistical analysis of the polyaffine deformations found. In the following, we first introduce our novel way of projecting SVFs onto the space of polyaffine transformations. Second, we formulate hierarchical polyaffine transformation trees and their estimation. Third, we conduct registration experiments on manual and hierarchical regions for femur and mandible bones. Fourth, we connect the clinical problem with our newly developed methodology by showing how it can be used for clustering in the context of population-based implant design.

2. Projection of Stationary Velocity Fields Onto the Linear Space of Log-Euclidean Polyaffine Transformations

In this section, we describe one of our main technical contributions, the projection of Stationary Velocity Fields (SVF) onto the space of Log-Euclidean Polyaffine Transformations (LEPT). This new projection allows us to estimate LEPTs by using the very efficient log-demons algorithm Vercauteren et al. (2009). This projection is independent of the registration problem. It can be considered as a new basis for SVFs, that can be closely linked to anatomical substructures (e.g. left and right side of the mandible bone). In Section 5, we show how it can be used in the context of registration.

In Arsigny et al. (2005), the authors introduced polyrigid and polyaffine transformations in the context of medical image analysis. Their works showed how to obtain diffeomorphic deformations by fusing transformations using ordinary differential equations (ODE). The estimation of the transformation parameters by numerically solving the ODEs was computationally expensive and therefore hard to use in practice. To tackle the computational burden the same authors introduced the fast Log-Euclidean polyaffine framework Arsigny et al. (2009). The authors redefined the problem by relying on the logarithms of the transformations. These logarithms are defined in the theory of Lie groups. In practice, this results in computations of these logarithms via matrix logarithms. The practical relevance of this theoretical reformulation of the problem was exploited in Commowick et al. (2008). In a first step, block-matching was used to estimate each affine transformation for a predefined region. In a second step, the algorithm fused the separately computed affine transformation using the Log-Euclidean polyaffine framework to ensure invertibility. The polyaffine transformation is not optimized but rather used as a regularization to combine affine transformations. In this

paper, we present a way to efficiently, directly and jointly estimate polyaffine transformation parameters for all regions in one step.

The novelty of our work is to take advantage of the efficiency of the log-demons algorithm to generate iterative solutions of SVFs, which we then project onto the space of LEPTs. In this way, we constrain the domain of possible solutions to the space of LEPTs. This computation scheme is detailed in Section 5. In this section, we focus only on the technicalities of the projections. To fully understand the projection we first define LEPTs (which includes the matrix logarithm and the exponential of matrices in terms of Lie group theory) and SVFs (which includes the exponential map of SVFs and its inverse). Finally, we have all ingredients to define our new projection.

2.1. Exponential and Logarithm of Matrices

Using homogeneous coordinates, the principal logarithm of affine transformations can be computed in a simple way. The main point here is that the principal logarithm of an affine transformation A_i is represented in homogeneous coordinates by the matrix logarithm of its representation M_i. This matrix logarithm takes the following form,

$$\log \begin{pmatrix} A_i & t_i \\ 0 & 1 \end{pmatrix} = \begin{bmatrix} M_i \\ 0 \end{bmatrix},$$

where log stands for the principal matrix logarithm, A_i is an 3×3 matrix and t_i an 3×1 vector, M_i is an 3×4 matrix and the index i can be ignored for now, it will become evident in Equation (2).

Using the principal logarithm of affine transformations one can associate to affine transformations a family of velocity vector fields in the following way,

$$v(x) = M_i \tilde{x}, \tag{1}$$

where $v(x)$ is an 3×1 and $\tilde{x}$ is an 4×1 vector.

The logarithms of affine transformations can be computed using matrix exponentials and the 'Scaling and Squaring' method as shown in Arsigny et al. (2009). For any square matrix, we have,

$$\exp \begin{pmatrix} M_i \\ 0 \end{pmatrix} = \exp \left(1/2^s \begin{bmatrix} M_i \\ 0 \end{bmatrix} \right)^{2^s}.$$

The key idea is that the matrix exponential is much simpler to compute for matrices close to zero. In this case, one can use only few terms of the infinite series of exponentials, since high-order terms will be completely negligible. An even better idea is to use Padé approximants, which provide excellent approximations by rational fractions of the exponential around zero with very few terms.

Exactly as for exponentials, we use 'Inverse Scaling and Squaring' method to compute matrix logarithms. Similarly, the idea is that logarithms are much simpler to compute for matrices close to the identity. To transform a matrix so that it is closer to the identity, the algorithm performs recursive computations of square roots, using,

$$\log \begin{pmatrix} A_i & t_i \\ 0 & 1 \end{pmatrix} = 2^s \log \left(\begin{bmatrix} A_i & t_i \\ 0 & 1 \end{bmatrix}^{1/2^s} \right).$$

Now let us introduce the LEPT for n regions with regions indexed by i. A linear combination of logarithms provides us with a SVF,

$$v_{\text{LEPT}}(x) = \sum_{i=1}^{n} w_i(x) M_i \tilde{x}, \tag{2}$$

where $w_i(x)$ are normalized weights for region i, i.e. $\forall x \in \Omega : \sum_{i=1}^{n} w_i(x) = 1$.

2.2. Exponential and Logarithm of Vector Fields

In Arsigny et al. (2006), the authors define the exponential $\exp(v)$ of a (smooth) vector field $v(x)$ as the flow at time 1 of the stationary ODE $\dot{x} = v(x)$. This generalizes the equivalence between one-parameter subgroups and exponential to SVFs v and diffeomorphisms ψ. This equivalence exists in the finite-dimensional case as described in the previous section. However, a proof for the existence and uniqueness of the logarithm $\log(\psi) = v$ is still an open research question. In this work we do not rely on it, since we only generate diffeomorphisms parameterized with SVFs using $\exp(v) = \psi$.

2.3. Our Contribution: Linear Projection

Given the velocity field v, we define the projection as the minimization of,

$$\text{proj}_M\, v = \underset{M}{\text{argmin}} \int_\Omega \lambda(x) ||v(x) - v_{\text{LEPT}}(x)||^2 dx = \underset{M}{\text{argmin}}\; C(M), \tag{3}$$

where $M = [M_1 \dots M_n]$, and λ is a binary mask indicating background voxels (if no mask is available $\forall x \in \Omega : \lambda(x) = 1$). This is a linear least squares problem. Using the Frobenius inner product, which generalizes the dot product to matrices, $||W||^2 = \text{Trace}(WW^\mathsf{T})$, we obtain the following directional derivative (see Appendix Appendix A on details of directional derivative of matrices):

$$\begin{aligned}
\partial_W C(M) &= \lim_{\epsilon \to 0} \frac{C(M + \epsilon W) - C(M)}{\epsilon} \\
&= \int_\Omega \lambda(x) \,\text{Trace}\left[\left(\sum_{i=1}^n w_i(x) W_i \tilde{x} \right) \left(\sum_{i=1}^n w_i(x) M_i \tilde{x} - v(x) \right)^\mathsf{T} \right] dx \\
&= \sum_{j=1}^n \text{Trace}\left[W_j \int_\Omega \lambda(x) w_j(x) \tilde{x} \left(\sum_{i=1}^n w_i(x) M_i \tilde{x} - v(x) \right)^\mathsf{T} dx \right],
\end{aligned}$$

where $w_i(x)$ are normalized weights, $\forall x \in \Omega : \sum_{i=1}^n w_i(x) = 1$. At the optimum, the directional derivative should be null in all directions W, i.e. $\partial_W C = 0$ for all matrices W. As we have $\text{Trace}(WA) = 0$ for all W iff $A = 0$, we end up with the system,

$$\sum_{j=1}^n M_i \Gamma_{ij} = B_i,$$

with

$$\Gamma_{ij} = \int_\Omega \lambda(x) w_i(x) w_j(x) \tilde{x} \tilde{x}^T dx, \tag{4}$$

$$B_i = \int_\Omega \lambda(x) w_i(x) v_c(x) \tilde{x}^T dx, \tag{5}$$

where Γ is symmetric and thus diagonalizable and the minimal norm solution is given by the pseudo inverse:

$$M = [M_1 \dots M_n] = B\Gamma^+. \tag{6}$$

2.4. New Metric on Stationary Velocity Fields

This new basis on SVFs gives us a new metric on SVFs,

$$\langle v_{\text{LEPT}_1}, v_{\text{LEPT}_2} \rangle_{L_2} = m_1^\mathsf{T} G m_2, \tag{7}$$

where $G = \Gamma \otimes I_3$, with Γ given by Equation (4) from the previous section. The symbol $\otimes$ denotes the Kronecker product, and I_3 represents the 3×3 identity matrix.

3. Hierarchical Structuring of Polyaffine Transformations in Trees

In this section, we assume that the weights of each region are already defined (the definition is deferred to the next section). The aim of this section is to build the tree structure.

Let us introduce the polyaffine transformation tree, which is a general formulation of Seiler et al. (2011b) for n regions and k levels. Let M_i^l be the ith 3×4 non null components of the matrix logarithm of affine transformations at level l and v^l be the LEPT (introduced in the previous section in Equation (2)) at level l:

$$v^l(x) = \sum_{i=1}^{n} w_i^l(x) M_i^l \tilde{x}, \tag{8}$$

with

$$\log \begin{pmatrix} A_i^l & t_i^l \\ 0 & 1 \end{pmatrix} = \begin{bmatrix} M_i^l \\ 0 \end{bmatrix}, \ \tilde{x} = \begin{bmatrix} x \\ 1 \end{bmatrix},$$

where

- $v^l(x)$ is a 3×1 vector at spatial position x (3×1 vector),
- $w_i^l(x)$ are scalar weights for regions i,
- A_i^l is the linear part (3×3 matrix) of the affine transformation,
- t_i^l is the translational part (3×1 vector) of the affine transformation.

In Fig. 3, the tree structure is illustrated for the first 3 levels on mandible CT data. Each node is assigned one spatial weight function $w_i^l(x)$ and one transformation M_i^l, and each node has two child nodes, if it is not a leaf node.

The tree structure builds the basis of our work, in the next two sections we elaborate on the definition of weight parameters $w_i^l(x)$ and the estimation of transformation parameters M_i^l.

4. Anatomy-Driven Definition of Regions

In this section, we describe the definition of weights for each regions using a hierarchy of Oriented Bounding Boxes (OBB) for two different cases: surface contours and voxel features from the image.

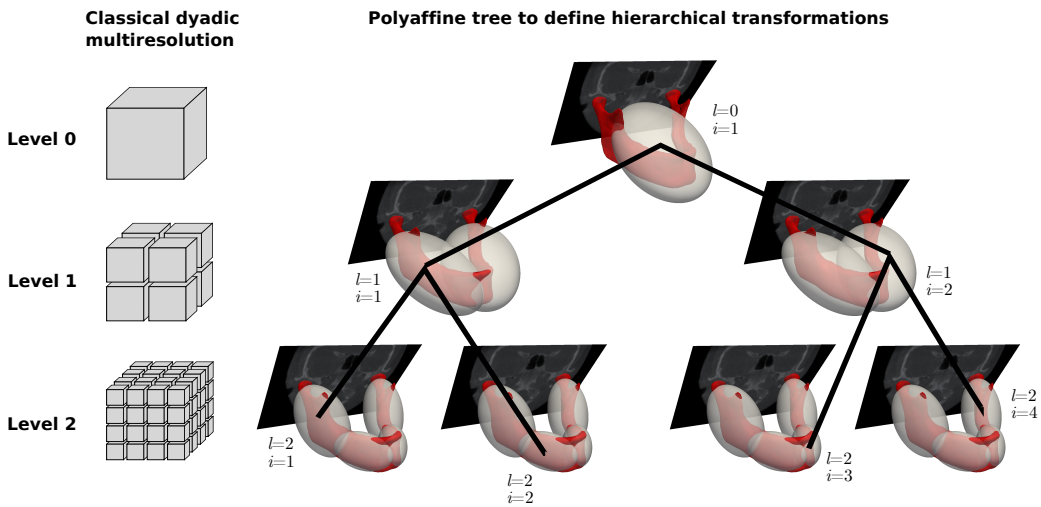

Figure 3: First three levels of space decomposition. **Left:** For comparison, a traditional approach using a dyadic multiresolution scheme. **Right:** Our new method decomposes the image domain using a tree of Gaussian weights. Each Gaussian weight represents one region and is visualized as one ellipsoids at σ. In addition to the contour (in red) extracted from the CT image, one CT slice is shown to stress that we work in the entire 3D image domain and not only with contours.

4.1. OBBTree on Contour Extracted from Template Image

The concept of OBB has been used extensively in computer graphics to speed up ray tracing and interference detection computations. In Gottschalk et al. (1996), the authors presented a hierarchical version and an algorithm to compute it efficiently. An OBBTree is a hierarchy of OBB's in 3D space. Let us first consider a surface (in our case a contour of the segmented CT image). The algorithm computes OBB's via Principle Component Analysis (PCA) of the vertex coordinates, which gives the orientation (principle component directions) and the extent (outmost projected point on the principle component). A refinement to avoid bias towards densely populated patches is to sample the convex hull of the vertex coordinates and approximate the analytic surface by a linear sum of all triangle areas. There are two ways of calculating the hierarchy, bottom-up and top-down. Top-down approaches start with all vertices and subdivide the points into two groups at every subsequent hierarchical level, whereas bottom-up approaches start by assigning one box per vertex and combine vertices until one box contains all vertices. We used the top-down approach, which divides the vertices into two groups

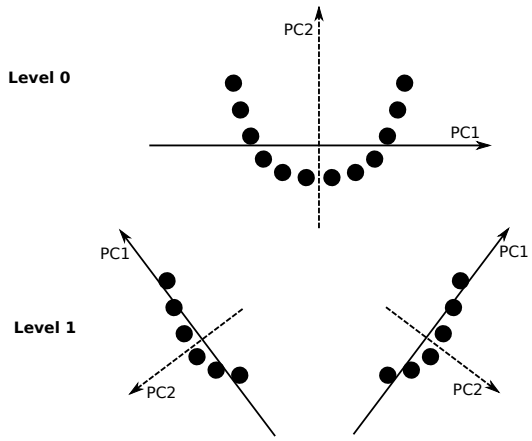

Figure 4: OBBTree algorithm applied to a toy example of points for level 0 and 1. At each level of the algorithm the pointset is split into two subsets along the first principle component (PC1) at the intersection with the second principle component (PC2).

by projecting the vertex coordinates onto the principle components, and uses the mean point as the group boundary. The algorithm stops once there are no more possible divisions along any component. The algorithm is illustrated in Fig. 4 for a toy example of a set of points. In our mandible application, points are replaced by vertices of the mandible contour.

4.2. OBBTree on Feature Image Extracted from Template Image

Instead of working on a subset of the image, e.g. the contour of mandible CT data as shown in the previous section, here we propose a method to use information from the entire image domain. There are many different ways to extract features from medical images and choosing the most appropriate one usually depends on the application and image modality at hand. In this work we use the following scalar feature image extracted from the template image $I_t(x)$,

$$\psi(x) = \log(1 + ||\nabla I_t(x)||^2), \tag{9}$$

where $\nabla I_t(x)$ is the gradient of the image. We take the logarithm of the gradient to be more robust against small changes in intensities due to noise. As a next step we generalize the PCA of vertices to feature weighted voxels. For this purpose we introduce the feature-weighted barycenter of region Ω_i^l

(discretized at voxel indices j),

$$\bar{x}_i^l = \frac{1}{\sum_{j \in \Omega_i^l} \psi(x_j)} \sum_{j \in \Omega_i^l} \psi(x_j) x_j, \tag{10}$$

where $\bar{x}_i^l$ is a 3×1 vector. The feature-weighted 3×3 covariance matrix of the region is,

$$\Sigma_i^l = \frac{1}{\sum_{j \in \Omega_i^l} \psi(x_j)} \sum_{j \in \Omega_i^l} \psi(x_j)(x_j - \bar{x}_i^l)(x_j - \bar{x}_i^l)^\mathsf{T}. \tag{11}$$

We perform singular value decomposition to obtain the principle components. Following the same strategy as in the original OBBTree algorithm, the region is split at the feature-weighted mean point orthogonal to the first principle component and the splitting procedure is recursively repeated in the two new created subregions.

Finally, the spatial weight functions $w_i^l(x)$ are defined as multi-dimensional Gaussian functions,

$$w_i^l(x) = (2\pi)^{-\frac{3}{2}} |\Sigma_i^l|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(x - \bar{x}_i^l)^\mathsf{T} (\Sigma_i^l)^{-1} (x - \bar{x}_i^l)\right), \tag{12}$$

with singular value decomposition,

$$\Sigma_i^l = V_i^l (\alpha S_i^l)^2 (V_i^l)^\mathsf{T}, \tag{13}$$

where

- V_i^l is a 3×3 rotation matrix describing the orientation,
- S_i^l is a 3×3 diagonal matrix of eigenvalues describing the scaling,
- α is a scalar representing the scaling parameter to control the variance in all three directions and thus the influence of regions and its neighbors. This parameter α is to be defined by the user and controls the smoothness of the deformation v^l.

OBBTree algorithms on contours are a special case of the presented feature image-based algorithm, when $\forall x \in C : \psi(x) = 1, \forall x \notin C : \psi(x) = 0$, where C is the mandible contour.

5. Estimation of Transformations

Up to now, we have defined a hierarchical tree of parametric deformations which is well fitted to the anatomy that we investigate. We now need a method to estimate these parameters. In this section, we describe the basics of the log-demons algorithm and how it is used to estimate polyaffine transformations. The main idea is to iteratively find the optimal polyaffine transformation tree that describes the correspondences found by the log-demons algorithm.

5.1. Log-Demons Algorithm

The log-demons algorithm finds diffeomorphic deformations to warp a moving image into a fixed image. The deformations are parameterized with Stationary Velocity Fields (SVF). A displacement vector field can be generated from a SVF through the exponential map. The exponential map $\exp(v)$ of a smooth SVF v is defined in Arsigny et al. (2006) as the flow at unit time, $\phi(x,1) = \exp(v(x))$, of the stationary ODE, $\partial\phi(x,t)/\partial t = v(\phi(x,t))$. To efficiently compute $\exp(v)$, Arsigny et al. (2006) proposed to use the scaling and squaring method. The theoretical motivation lies in the generalization of Lie Group theory to the infinite dimensional case. Although there are still some open questions to be resolved, in practice this approach provides good results. One interesting point of this registration framework is the efficient optimization in the domain of SVFs. This property explains the denomination of log-domain (or simply log-demons) registration.

The general form of the demons algorithm was formulated in Cachier et al. (2003) and later in Vercauteren et al. (2009) an implementation with SVFs instead of displacement vector fields was presented. The goal is to find v^l that warps the moving image I_s into the fixed image I_t (or resamples I_s in I_t), by minimizing the cost functional,

$$C\left(I_t, I_s, v_c, v^l\right) = \sigma_i^{-2}\,\mathrm{Sim}\left(I_t, I_s, v_c\right) + \sigma_x^{-2}\,\mathrm{dist}\left(v_c, v^l\right)^2 + \sigma_T^{-2}\,\mathrm{Reg}\left(v^l\right),$$

where

- Sim is the similarity between two images, we use the sum of squared distances on the image intensities, $\mathrm{Sim} = \frac{1}{2}||I_t - I_s \circ \exp(v_c)||^2$,
- dist is the hidden term linking Sim and Reg, we use $\mathrm{dist} = ||v_c - v^l||$,
- Reg is the regularization term, we use $\mathrm{Reg} = \left|\left|\nabla v^l\right|\right|^2$.

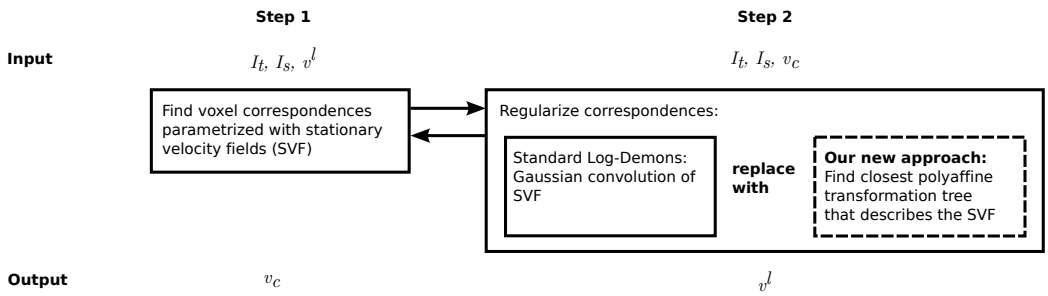

Figure 5: A schematic overview of the log-demons algorithm. The log-demons consists of two major steps: correspondence finding and regularization. The polyaffine transformation trees are estimated in the regularization step. The input and output row show which variable is optimized and which variable is fixed during each step.

Each term has a weighting parameter σ_i, σ_x and σ_T. As shown in Vercauteren et al. (2009), $C\left(I_t, I_s, v_c, v^l\right)$ can be optimized alternatively over the variables v_c, the correspondence velocity field computed by the first optimization part of the log-demons algorithm, and v^l. First, minimizing,

$$C_{\text{sim}}(I_t, I_s, v_c, v^l) = \sigma_i^{-2}\,\text{Sim}\left(I_t, I_s, v_c\right) + \sigma_x^{-2}\,\text{dist}\left(v_c, v^l\right)^2,$$

with respect to v_c and fixed v^l. Second, a Gaussian smoothing of v_c, which minimizes,

$$C_{\text{reg}}\left(I_t, I_s, v_c, v^l\right) = \sigma_x^{-2}\,\text{dist}\left(v_c, v^l\right)^2 + \sigma_T^{-2}\,\text{Reg}\left(v^l\right),$$

when the regularization term is $\text{Reg}\left(v^l\right) = \left|\left|\nabla v^l\right|\right|^2$. Alternating between these two steps, as shown in Fig. 5, allows for a very efficient registration.

5.2. Integration of the Estimation into the Log-Demons

To optimize the polyaffine transformation tree, we remove the regularization term but we constrain the velocity field v^l to be issued from a polyaffine transformation, i.e. to have the form specified in Section 3, Equation (8). First, we minimize the functional with respect to v_c,

$$\begin{aligned} C(I_t, I_s, v_c, v^l) &= \sigma_i^{-2}\int_\Omega (I_t - I_s \circ \exp(v_c))(x)^2 dx + \\ &\quad \sigma_x^{-2}\int_\Omega \left|\left|v_c(x) - v^l(x)\right|\right|^2 dx, \end{aligned} \tag{14}$$

while the first term describes the sum of squared differences (SSD) image metric, the second term is called the hidden term that allows for the splitting of the optimization. What differs from the log-demons algorithm introduced in Vercauteren et al. (2009) is that we replace the regularization step by a projection to find the closest polyaffine transformation tree. This requires only a modification in the second step of the optimization, schematically shown in Fig. 5.

Now given v_c, we can solve for M_i^l using linear least squares. Here we use our newly introduced projection of SVFs onto LEPTs as described in Section 2.3,

$$\text{proj}_{M^l} v_c = \underset{M^l}{\text{argmin}} \int_\Omega \lambda(x)||v_c(x) - \sum_{i=1}^{n} w_i^l(x) M_i^l \tilde{x}||^2 dx.$$

As shown in Equation (6), we end up with the following linear system of equation to solve,

$$M^l = [M_1^l \dots M_n^l] = B\Gamma^+.$$

Here, the matrix Γ does not change for different velocity field observations as long as the mask image and weights are fixed, hence only B needs to be recomputed for each iterative step of the log-demons optimization. The sequential per-level estimation with the log-demons algorithm is shown in Algorithm 1.

Algorithm 1 Estimation of Polyaffine Transformation Tree

Sequentially estimate levels $l = 0, \dots, k$

- Initialize demons with previous level $v^l = v^{l-1}$ (for starting level $v^0 = 0$)
- Precompute Γ
- Iterate until convergence
  - Compute correspondence SVF v_c
  - Compute residual SVF $v_r = v_c - v^{l-1}$
  - Project v_r onto $M_1^l, \dots, M_n^l$ by solving the linear least square problem $M^l = B\Gamma^+$

6. Femur Bone Registration with Manual Regions

In this section, we present validation results of our registration for manually defined regions. We first present the clinical application and then show how our registration performs w.r.t. manual landmark measurements conducted by two medical experts. The results show that the polyaffine method improves the accuracy of the standard log-demons registration.

6.1. Clinical Problem

Tumor excision is the primary treatment of aggressive or recurrent benign bone tumors and malignant bone sarcomas. This requires an invasive surgical intervention that entails a residual bone defect, which can be reconstructed with a fresh frozen bone allotransplantation. In orthopaedics, recent improvements in diagnostic and therapeutic techniques have produced an increase of the patient survival as well as a reduction of the complication rate Muscolo et al. (2005). It has been shown that the selection of bone allograft in terms of shape and size is crucial to prevent changes in joint mobility and load distribution, which can lead to joint fractures and early joint degeneration Enneking and Campanacci (2001). However, current selection approaches are very time-consuming, mostly based on manual measurements performed directly on the bones or on three-dimensional models reconstructed from images.

We present a method to perform allograft selection through image registration. We focus on six landmarks extracted from the velocity field obtained through registrations. The six landmarks define three distances on the bone (Fig. 7): Transepicondyle distance (A), anterior-posterior distance in the medial condyle (B) and in the lateral condyle (C). The clinical aspects of this work were recently published in Ritacco et al. (2012). Here we use the results to show the validity of our registration compared to measurements performed by two medical experts on the landmarks defining A, B and C.

6.2. Special Case of a Polyaffine Transformation Tree

To handle rotational misalignments of the femoral head and condyles we propose to split the bone into three regions. This results in a tree with two levels as shown in Fig. 6. Level 0, represents a global affine registration. Level 1, represents the division into three parts. The middle section of the femur is defined as shaft region. The distal and the proximal end are condyle and head region, respectively. We initialize the tree with an anisotropic scale transform computed from the enclosing bounding box of the entire femur.

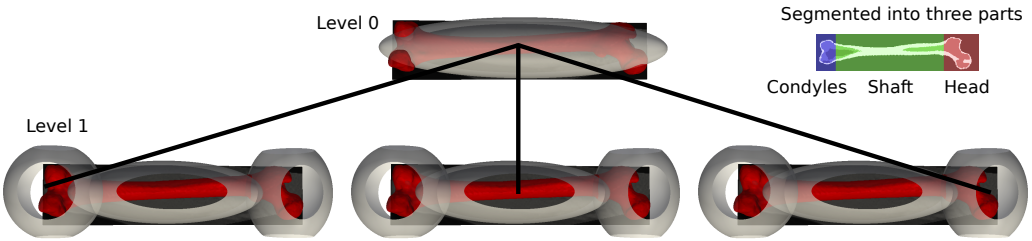

Figure 6: Ellipsoids representing Gaussian weights at σ for levels 0 and 1 (at the image boundary ellipsoids are cut). The Gaussian weights are defined through the labeled image shown on the top right, which splits the femur into three parts: Condyles, shaft and femoral head. The red femur is the surface extracted from the template image.

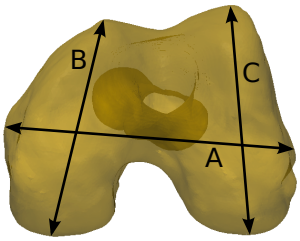

Figure 7: Six femur landmarks that were measured by medical experts, which define the three distances A, B and C.

6.3. Validation with Landmark Measurements by Medical Experts

We register 146 CT images of femur bones with voxel size 1 mm. The ground truth of the ABC measurements was obtained by manually selecting the landmarks on 3D reconstructed surface models. Two medical doctors performed these measurements independently and one expert repeated the measurements after a few days to yield the intraobserver variability. We compare four registration methods: First, an anisotropic scaling computed from the enclosing bounding boxes, referenced hereafter with *Scale*. Second, a standard log-demons, referenced hereafter with *LogDemons*. Third, the newly developed 2-level polyaffine tree, referenced hereafter with *Tree*. Forth, a standard log-demons initialized with the transformation obtained from the *Tree*, referenced hereafter with *Tree relaxed*.

The results are shown in Fig. 8. For landmark distance A, the *Tree* is more accurate than the *LogDemons*. This is quite surprising, because

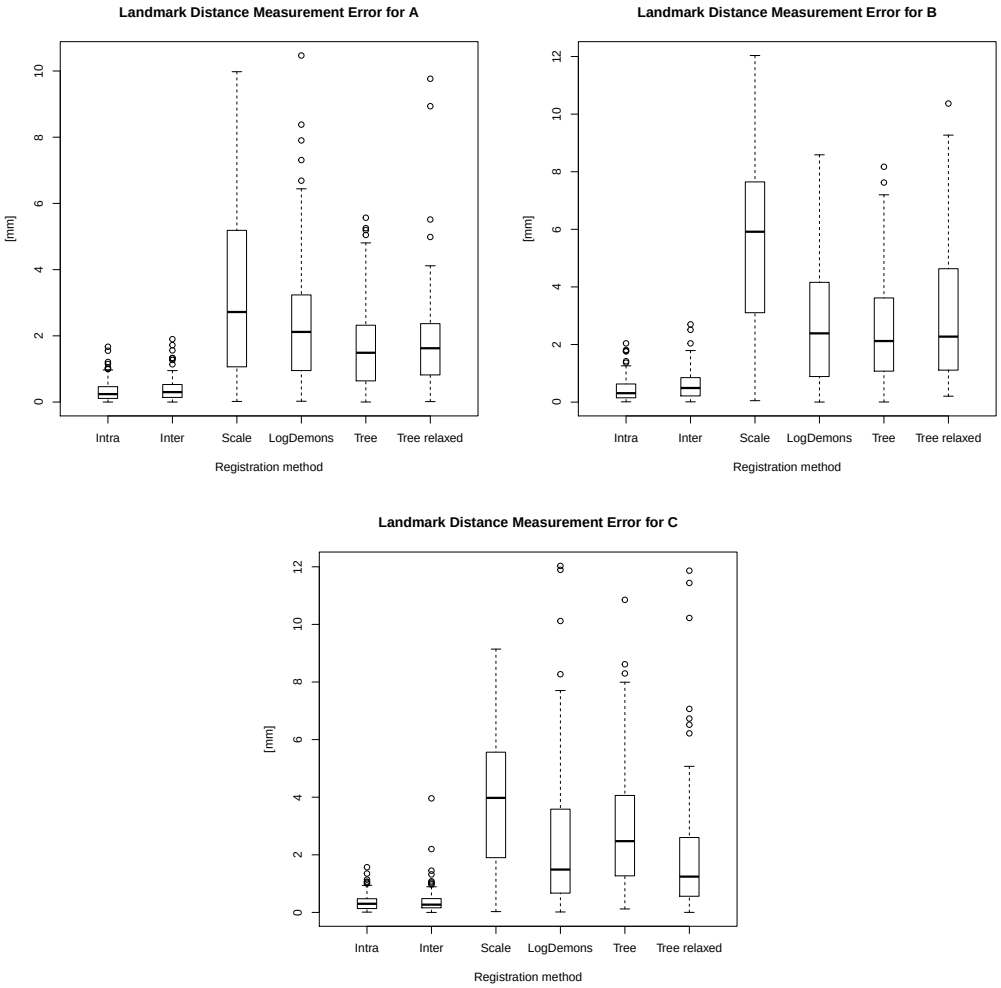

Figure 8: Boxplots of measurement errors (circles represent outliers) computed for each registration method and compared to *Intra* and *Inter* rater error conducted by medical experts. *Scale* represents a global scale transform; *LogDemons* a standard log-demons with scale initialization; *Tree* a 2-level tree with manual regions; *Tree relaxed* a standard log-demons with Tree initialization.

the number of parameters needed to describe the *LogDemons* are three per voxel, with an image size of $120 \times 138 \times 542$, this results in a total of 10^7 parameters, whereas for the *Tree* transformation we have one region at level 0 (12 parameters) and three regions at level 1 (36 parameters), resulting in a total of 48 parameters. *Tree relaxed* shows two outliers not present in *Tree*, indicating the robustness of estimating only a few parameters as opposed to an entire field. For landmark distance B, *LogDemons* and *Tree* perform similarly as for distance A, and *Tree relaxed* shows one outlier not present in *Tree*. For landmark distance C, *Tree* shows a higher median than *LogDemons* and *Tree relaxed*. Here *Tree relaxed* performs favorably.

For all three landmarks, we can state that either the *Tree* or *Tree relaxed* are more accurate and/or more robust towards outliers. The results for *Tree* are particularly impressive, since the number of parameters needed to describe the geometrical deformation is five orders of magnitude smaller. The reduction in the number of parameters to estimate during the registration could explain why the results are in all cases more robust. We believe that these improvements stem from considering torsion and rotation in head and condyles which is implicitly modeled by dividing the bone into three parts and is not considered in a standard non-linear registration. We can thus incorporate our prior knowledge about the anatomy into the registration process by simply defining a rough mask of expected regions. This clearly indicates the power of the new method to capture the anatomical variability.

7. Mandible Bone Registration with Hierarchical Regions

In the previous section, we showed registration results for manual regions defined on femur bones. This is a special case within the family of polyaffine transformation trees, where the tree has two levels (level 0, one region; level 1, three manually defined regions). Now we show the more general case by computing a hierarchy of regions. We show its usefulness for mandible bones, where we encounter different variability at different levels. In this kind of anatomical setting, manual regions are not tractable anymore, due to the increasing number of regions with number of levels (grows quadratically in the case of binary trees).

7.1. Weights Derived from Oriented Bounding Boxes

The multiscale and hierarchical regions are computed using the weighted OBBTree algorithm on feature images (as described in Section 4.2). The

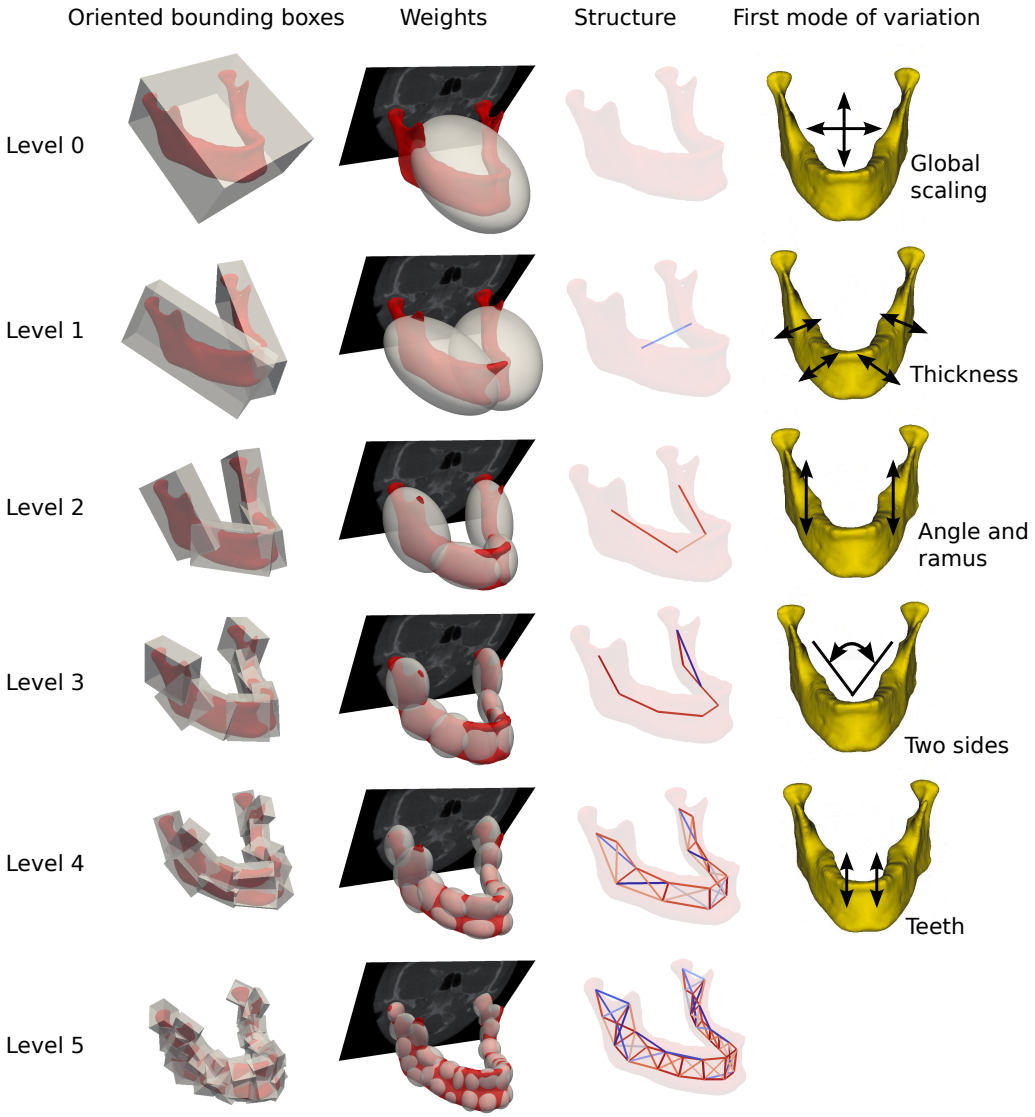

Figure 9: The red mandible is the surface extracted from the template image. **Column 1:** Oriented bounding boxes computed using the algorithm presented in Section 4. The Gaussian weights are driven by the gradient in the CT image. **Column 2:** Ellipsoids representing Gaussian weights at σ for levels 0 to 5. The parameters of the Gaussian weights are derived from the OBB. **Column 3:** Structure imposed by the weights. The color of the edges encode correlations between regions, ranging from low (blue=0.4) to high (red=1). **Column 4:** First PCA mode at each level, showing the main residual variation at that level.

resulting boxes are visualized in Fig. 9 (first column). Starting at level 1, we observe the division of the mandible into left and right side, followed by subdivisions into clinical regions as described in Fig. 1, to tooth-sized regions, and finally even finer scales (not visualized here), which are harder to interpret clinically. From the boxes, we computed the Gaussian weights with a scaling parameter of $\alpha = 1$. The Gaussian weights visualized as ellipsoids are shown in Fig. 9 (second column). In the third column, the correlation of weights are depicted according to,

$$\mathrm{Corr}_{ij} = \Sigma_{ii}^{-1/2} \Sigma_{ij} \Sigma_{jj}^{-1/2}, \tag{15}$$

where Σ_{ij} is defined in Equation (11), and Corr_{ij} is decomposed using singular value decomposition to extract the major axis of correlation. The correlation of the major axis is colored coded from cool (blue=0.4) to warm (red=1). The graph structure clearly reveals the intrinsic underlying dimensionality of the object at each scale, going from a curve to a ribbon and finally in some areas locally to a 3D volume. It would be interesting to study if this could be a robust alternative to the medial axis or surface representation.

7.2. Estimation of the Transformation Tree

Rigid registration is considered as a pre-processing step to the work presented here. Indeed, intensity-based rigid registration procedures failed due to the angle between left and right side of the mandible, causing the registration to either fit the left or the right side. Instead, we extract surfaces from segmented images and align them according to the principle directions of their vertices. For the population-based analysis of mandible CT images, we register all 43 CT images to a template and analyze the deformations obtained. The template (35 year old male) was selected through visual examination. For the implant design study presented in Section 8, we need registrations up to level 4. In this section, we analyze all levels and compare them to two other types of registration.

In Fig. 10, the mean squared error (MSE) of intensity difference (fiducial localization error), computed over the union of template mask and subject mask, for three different types of registrations are shown.

- Type 1 (*S0* and *S0R*): Standard log-demons *S0R* with a global affine initialization computed using a tree with one level containing the root node *S0*.

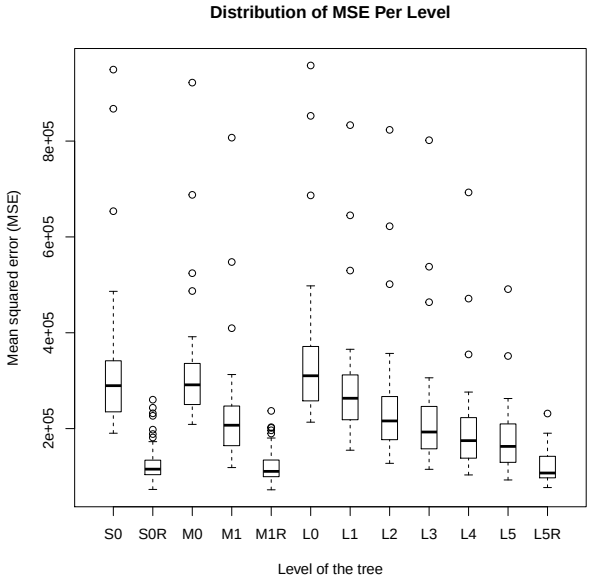

Figure 10: Performance of the multiscale polyaffine registration. Mean squared error (MSE) of intensity difference (fiducial localization error) calculated on the union of template and subject mask image for each of the 43 registrations (without the template to itself registration). Explanation of boxplots (circles represent outliers) starting from the left hand side: *S0:* Level 0 of manual tree defined through mask image. *S0R:* Standard log-demons initialized with S0. *M0:* Level 0 of manual tree defined through mask image. The results differ from S0 because the normalization of weights is computed over all levels. *M1:* Level 1 of manual tree defined through label image. *M1R:* Standard log-demons initialized with M1. *L0 to L5:* Level 0 to level 5 of the data-driven multiscale and hierarchy of regions. *L5R:* Standard log-demons initialized with L5.

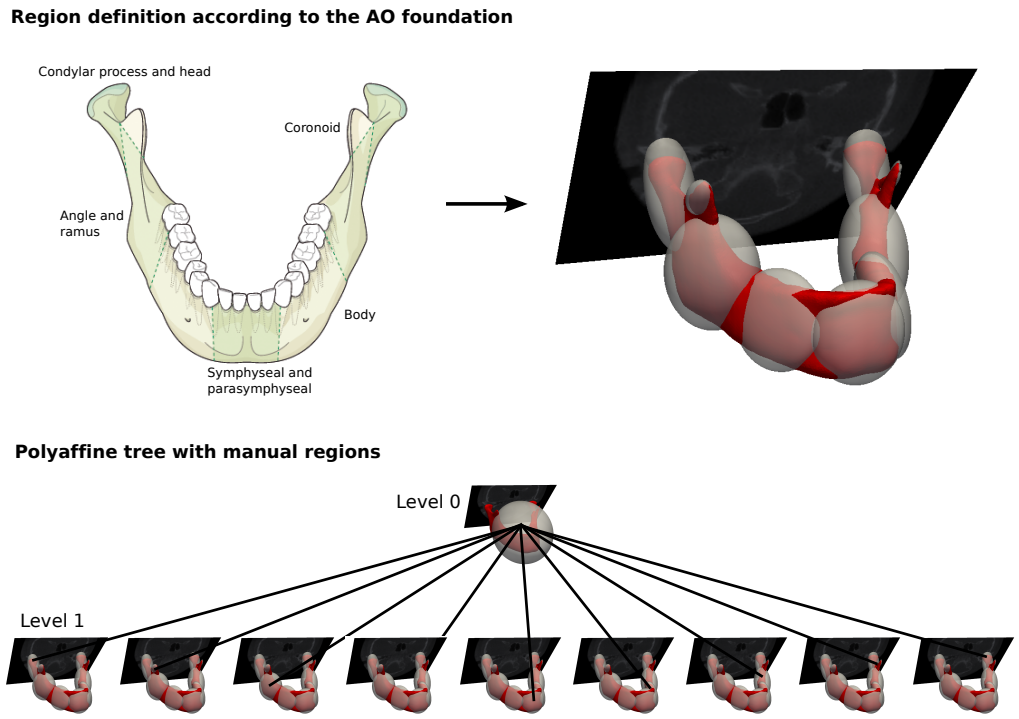

Figure 11: Manual regions defined on the reference image according to the AO foundation classification scheme.

- Type 2 (*M0*, *M1* and *M1R*): A manually defined tree of regions as depicted in Fig. 11. At level 0, one region to capture the global affine variability *M0*, and at level 1, 9 regions defined according to the AO foundation classification scheme *M1*. The transformation parameters are estimated using Section 5. Standard log-demons *M1R* initialized with *M1*.

- Type 3 (*L1–L5* and *L5R*): Data-driven multiscale and hierarchical weight tree derived with the algorithm presented in Section 4 *L1–L5*. The estimation of transformation parameters uses Section 5. Standard log-demons *L5R* initialized with *L5*.

Figure 10 shows two main results:

- Decrease of MSE with increasing amount of regions (levels). The MSE converges, which provides evidence of the consistency of our method

(although the finer levels are not statistically analyzed in this work).

- All three relaxed registrations *S0R*, *M1R* and *L5R* are in the same range, with *L5R* having less outliers. This suggest a more robust registration due to the stepwise initialization across levels.

Unfortunately, we cannot evaluate the accuracy here as we do not have manual measurements provided by medical experts for mandibles. In the future, we plan to validate the accuracy in a similar fashion as described in the previous section for femurs.

7.3. Analysis of Obtained Tree Transformations

One important aspect of our tree structure lies in its power of decomposing features into different scales, this can be shown by performing a hierarchical PCA for mandibles. The per level PCA (Fig. 9, column 4) can be interpreted as follows: (Level 0) global scaling; (Level 1) thickness; (Level 2) reorientation in the region of the masseter; (Level 3) relative displacement of condyles and coronoid processes; (Level 4) change in teeth region. This gives a visual validation of the usefulness of per level hierarchical statistical analysis, clearly distributing features across different scales.

We encountered instability starting at level 5, where the PCA modes become anatomically unrealistic. We are currently working on a Bayesian framework to make the registration more robust for higher levels.

8. Implant Design Based on Statistical Analysis of Transformations

In this section, we show how our registration method can be used for implant design. In order to analyze the groupwise deformations modeling the different anatomies in our population, we first draw the connection between the methodology and the clinical problem of population-based implant design in Section 8.1 (depicted in the first and second row of Fig. 12). Then, we perform (Section 8.2) a k-means clustering of transformations for selected regions (results depicted in the last row of Fig. 12). Each mean of each cluster represents a specific mandible geometry (surface and voxel intensities), which can be used to drive the design of an implant.

Our dataset consists of 43 mandible CT images of healthy patients aged between 23 to 88, with a median age of 65.5 and four patients with unknown age; 16 patients were female and 26 male. The images are resampled during rigid registration from an original spacing of 0.4 mm (uniform) to a more computational tractable spacing of 1.25 mm.

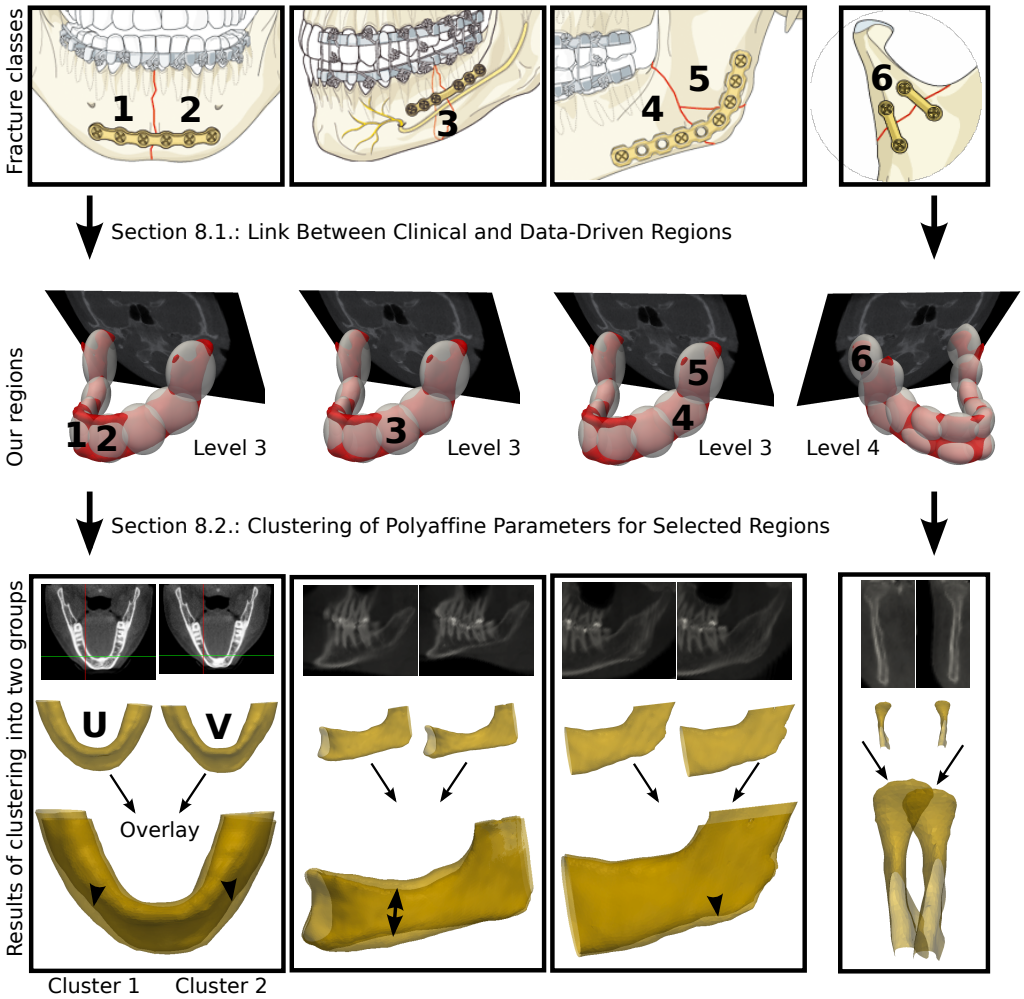

Figure 12: Overview of the entire work flow. Starting with the fracture classes given by the AO foundation, we identify 6 regions that are close to fracture sites. We then identify corresponding regions that were retrieved during our new registration method. The three columns on the left are regions identified on level 3, the one on the right is one region at level 4. K-means clustering for each column is performed. The last row, shows the computed clusters: One slice extracted form the two mean images of each cluster, surface of the mean image of each cluster and overlay of the two mean surfaces for comparison.

8.1. Link Between Clinical and Data-Driven Regions

Given the regions that were found on the template image in Fig. 9, we identify the corresponding clinically motivated regions from Fig. 1. This allows us to identity a subset of clinically important parameters. We can perform a statistical analysis for each implant design study on 12 (one region) or 24 (two regions) parameters. At this point we do not enforce a perfect overlap between the data-driven regions and the clinical regions. Work in progress focuses on an approach to incorporated prior knowledge into the tree to enforce a perfect overlap for certain regions, while allowing for data-driven regions in areas where there is no anatomical information available (i.e. very fine structures at fine levels). The link between automatic OBB-Tree and anatomical regions is shown in Fig. 12. For "symphyseal and parasymphyseal", "body", and "angle and ramus" we select regions at level 3; for "condylar process and head" one region at level 4 is identified.

8.2. Clustering of Polyaffine Parameters for Selected Regions

The link between clinical and data-driven regions, presented in the previous section, allows us to focus on each part of the bone independently. For each part the goal is to find clusters representing distinguishable surface and intensity features. We apply k-means for the clustering of each part. The selection of the number of clusters, is a crucial decision in k-means, and in our application it is up to the implant designer to select how many different implants to build as each additional implant causes additional costs. In Fig. 13, the reduction of Within groups Sum of Squares (WSS) as a function of the number of clusters indicates a gradual decrease. In our experiments we use two clusters. In Fig. 12 on the last row, the different surface shape and intensity features are shown. To illustrate the difficulty of finding these feature differences by visual inspection, Fig. 14 shows one slice of each mandible in the transversal plane. Each slice represents the transformed template slice at level 3. As one can see it is hard to distinguish the shape and intensity feature by pure inspection of individual images.

9. Discussion and Conclusion

In this work, we proposed a new approach to population-based implant design for mandibles. We presented a new method that is able to link the registration directly to the clinical relevant anatomical regions. Experiments on 146 CT images of femur bones showed that our approach (with manual

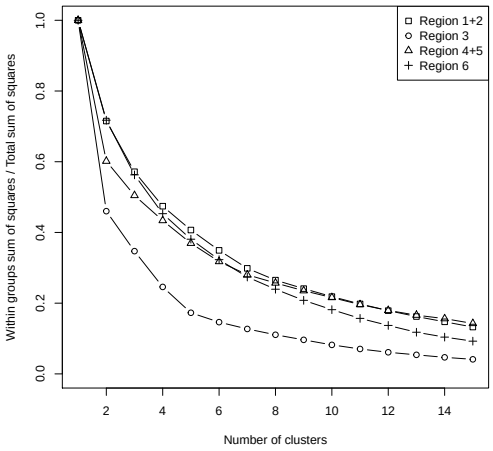

Figure 13: K-means clustering results for all four regions. Region 3 shows a decrease by approximately half of the WSS at two clusters.

regions) is more accurate than the standard log-demons. The ground truth was provided by two medical experts, who manually identified three landmark distances. Furthermore, we illustrated the usefulness of the link between registration and clinical regions (data-driven hierarchical regions) through a clustering of 43 mandibles CT images, considering not only the surface but also the volumetric information.

In contrast to current population-based implant design methods we could show a direct connection between clinically relevant regions and our methodological framework. Our findings are in accordance with the clinical literature. The authors in Watanabe et al. (2010) presented a morphological study on manual measurements performed on 79 Japanese patients on "symphyseal and parasymphyseal", "body" and "angle" regions. A large variability in terms of standard deviation ranging from 3.5 mm to 3.8 mm in height and 1.3 mm to 3.2 mm in width was reported, supporting the evidence of high variability in these regions that were reflected in our results. In addition to these surface morphological findings, the volumetric shape informations could now be used for biomechanical simulations on optimal screw placement by taking advantage of the correlation between image intensities and material properties, along the line of McBroom et al. (1985). This implies an addi-

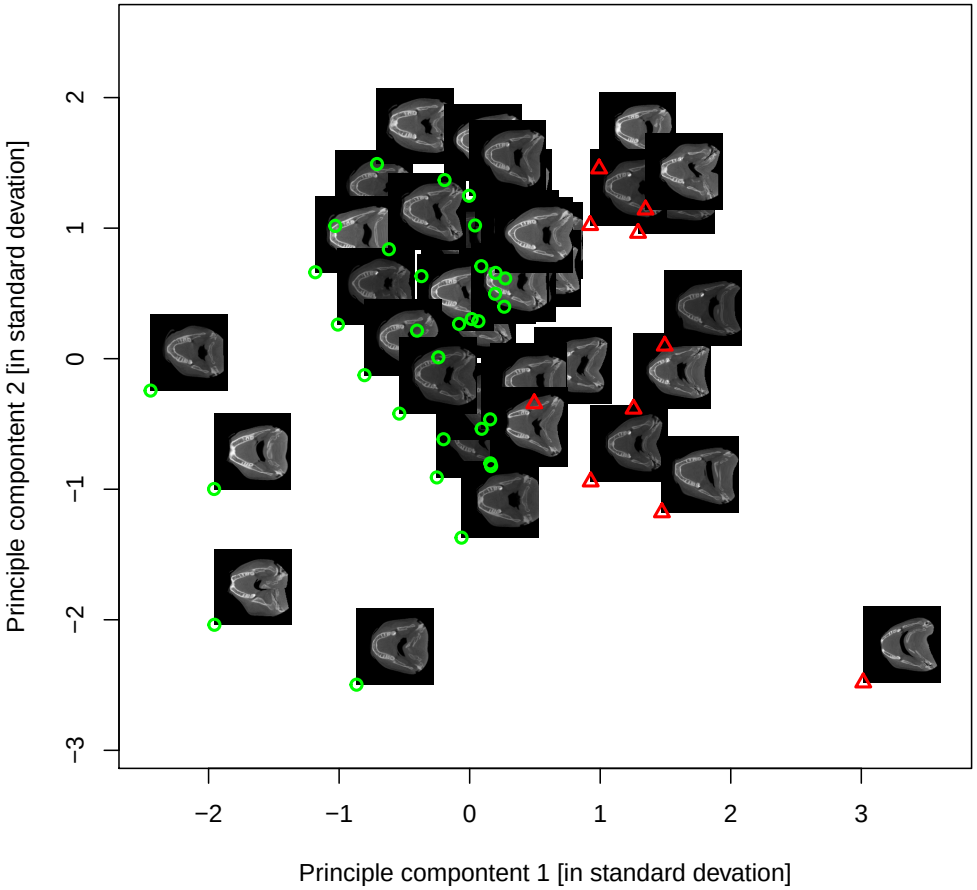

Figure 14: K-means clustering results for "symphyseal and parasymphyseal" (region 1+2). Red triangles represents the V-shaped mandibles; green circles represent the U-shaped mandibles. The slices are generated by warping the template slice with transformations at level 3. The 24-dimensional data points (two affine transformations) are projected onto the two largest principle components for visualization. The clustering is done in the 24-dimensional space.

tional statistical analysis of the intensities in the template space. Preliminary results in this direction were shown in Bonaretti et al. (2011).

Our method depends on three major parameters: First, the selection of the template image I_t, which serves as the basis of the weight definition, and thus influences not only the way transformations are estimated but also which regions are selected for clinical analysis. Hence, by changing the template we expect to obtain slightly different results. Due to the fact that weights can be easily visualized, an assessment of the quality of the selected template can be done by clinicians quite reliably. To further improve the robustness of this procedure, we plan to introducing a groupwise weight definition framework, in which we search for the optimal configuration of weights considering all images available. Second, the scaling parameter α of the covariance matrix, which controls the smoothness of the final deformation. In the current work, this parameter was set heuristically. An alternative approach would be to find the optimal scaling parameter by evolving the regions (and thus weights) during the transformation estimation phase. Third, the total number of levels k to consider for registration. In the presented clinical application this parameter is set to level 3 and 4 given the clinical regions that are of interest. For other applications, it might be useful to define a stop criterion based on the residual variability. One step further, would be to enforce a stop criterion per node instead of the entire level. This would further reduce the number of parameters used to describe a deformation and improve the direct analysis of the registration results by clinicians.

Our methods can be decomposed into three major building blocks: First, the structure of the polyaffine trees. At the moment this is a binary tree with a fixed number of levels and fixed number of nodes. The tree structure with its number of levels and number of nodes could be relaxed and adapted by learning an optimal structure from a population of images. This optimization could be accomplished in a groupwise fashion by imposing an additional penalty favoring certain structure over others. Second, the definition of regions. By reformulating the OBBTree subdivision process in a more probabilistic framework, e.g. Gaussian Mixture Models, we could introduce priors, e.g. independence between regions on the polyaffine transformation trees, leading to an even lower dimensional representation. We can see at this point that the first two building blocks are very strongly linked. Third, the estimation of transformations. In the current implementation, we use the sum of squared differences as our similarity measures in the log-demons algorithm. This could be extended to consider other similarity measures.

Furthermore, other methods than the log-demons algorithm could be used to drive the estimation. All methods that are parametrized with SVFs are potential candidates. Additionally, by changing the coarse to fine transformation estimation to a joint estimation over all levels, we believe we can further improve the results in terms of robustness toward outliers (by taking care of artifacts at finer levels) and improving interpretability (by skipping levels). At the core of our new method lays the introduction of a new basis for SVFs. As we showed in our experiments, this basis has very close links to anatomical substructures. In the future, this method has the potential to discover the hidden and possibly sparse structure of the anatomy. We currently work on a method that imposes sparseness on the basis on a groupwise level by formulating our method in a Bayesian framework.

Acknowledgements

We would like to thank Prof. Tateyuki Iizuka, Department of Cranio-Maxillofacial Surgery, University of Bern, Switzerland, for providing us with valuable clinical feedback. We thank Dr. Lucas Ritacco from the Department of Medical Informatics and Institute of Orthopedics Carlos E. Ottolenghi, and Italian Hospital of Buenos Aires in Argentina, and Michael Kistler from the University of Bern, for manually measuring landmarks on 146 femur CT images. This work was funded by the Swiss National Science Foundation (SNSF). The collaboration between Switzerland and France is partially supported by the Rectors' Conference of Swiss Universities (CRUS).

Appendix A. Directional Derivative of Matrices

We use the method of directional derivatives to derive C w.r.t. M in the direction of W in Section 2.3. When the directional derivative is zero we have the minimal solution (in a least square sense) for M.

We present three definitions and properties that we later use for the method. First, we use the following definitions for directional derivative:

$$\partial_W C(M) = \frac{\lim_{\epsilon \to 0}(C(M + \epsilon W) - C(M))}{\epsilon} \iff$$
$$C(M + \epsilon W) = C(M) + \epsilon \partial_W C(M) + \mathcal{O}(\epsilon^2).$$

Second, we use the Frobenius inner product to write,

$$\partial_W C(M) = \left\langle W, \frac{\partial C(M)}{\partial M} \right\rangle = \text{Trace}\left(W^T \frac{\partial C(M)}{\partial M} \right).$$

Third, we use the following properties of the trace,

$$\begin{aligned}\mathrm{Trace}(W(\nabla C)) &= \mathrm{Trace}((\nabla C)W) = \mathrm{Trace}(W^T(\nabla C)^T),\\ \mathrm{Trace}(W + (\nabla C)) &= \mathrm{Trace}(W) + \mathrm{Trace}(\nabla C).\end{aligned}$$

Now, the method is to compute the Taylor expansion of $C(M + \epsilon W)$, then we take the non-W part of the first order term and by using the properties of the trace, we rewrite it as,

$$\partial_W C(M) = \mathrm{Trace}(W^T(\nabla C)), \tag{A.1}$$

now by identification, we have,

$$\frac{\partial C(M)}{\partial M} = \nabla C.$$

References

Arsigny, V., Commowick, O., Ayache, N., Pennec, X., 2009. A Fast and Log-Euclidean Polyaffine Framework for Locally Linear Registration. Journal of Mathematical Imaging and Vision 33, 222–238.

Arsigny, V., Commowick, O., Pennec, X., Ayache, N., 2006. A Log-Euclidean Framework for Statistics on Diffeomorphisms, in: Larsen, R., Nielsen, M., Sporring, J. (Eds.), MICCAI 2006, Springer, Heidelberg. pp. 924–931.

Arsigny, V., Pennec, X., Ayache, N., 2005. Polyrigid and Polyaffine Transformations: A Novel Geometrical Tool to Deal with Non-rigid Deformations – Application to the Registration of Histological Slices. Medical Image Analysis 9, 507–523.

Bonaretti, S., Seiler, C., Boichon, C., Büchler, P., Reyes, M., 2011. Mesh-Based vs. Image-Based Statistical Model of Appearance of the Human Femur: A Preliminary Comparison Study for the Creation of Finite Element Meshes, in: MeshMed, MICCAI Workshop.

Bou-Sleiman, H., Ritacco, L.E., Nolte, L.P.P., Reyes, M., 2011. Minimization of Intra-Operative Shaping of Orthopaedic Fixation Plates: A Population-Based Design, in: Peters, T., Fichtinger, G., Martel, A. (Eds.), MICCAI 2011, Part II, Springer, Heidelberg. pp. 409–416.

Bou-Sleiman, H., Seiler, C., Iizuka, T., Nolte, L.P., Reyes, M., 2012. Population-Based Design of Mandibular Plates Based on Bone Quality and Morphology, in: Nicholas, A., Hervé, D., Polina, G., Kensaku, M. (Eds.), MICCAI 2012, Springer, Heidelberg.

Buerger, C., Schaeffter, T., King, A.P., 2011. Hierarchical Adaptive Local Affine Registration for Fast and Robust Respiratory Motion Estimation. Medical Image Analysis 15, 551–564.

Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N., 2003. Iconic Feature Based Nonrigid Registration: The PASHA Algorithm. Computer Vision and Image Understanding 89, 272–298.

Cervantes, T.M., Slocum, A.H., Seldin, E.B., 2012. Design and Experimental Evaluation of Adjustable Bone Plates for Mandibular Fracture Fixation. Journal of Biomechanics 45, 172–178.

Commowick, O., Arsigny, V., Isambert, A., Costa, J., Dhermain, F., Bidault, F., Bondiau, P.Y., Ayache, N., Malandain, G., 2008. An Efficient Locally Affine Framework for the Smooth Registration of Anatomical Structures. Medical Image Analysis 12, 427–441.

Craig, M., Bir, C., Viano, D., Tashman, S., 2008. Biomechanical Response of the Human Mandible to Impacts of the Chin. Journal of Biomechanics 41, 2972–2980.

Ellis, E., Moos, K.F., El-Attar, A., 1985. Ten Years of Mandibular Fractures: An Analysis of 2,137 Cases. Oral Surgery, Oral Medicine, and Oral Pathology 59, 120–129.

Enneking, W.F., Campanacci, D.A., 2001. Retrieved Human Allografts: A Clinicopathological Study. J Bone Joint Surg Am 83, 971–986.

Fedok, F.G., Van Kooten, D.W., DeJoseph, L.M., McGinn, J.D., Sobota, B., Levin, R.J., Jacobs, C.R., 1998. Plating Techniques and Plate Orientation in Repair of Mandibular Angle Fractures: An in Vitro Study. The Laryngoscope 108, 1218–1224.

Freiman, M., Voss, S.D., Warfield, S.K., 2011. Demons Registration With Local Affine Adaptive Regularization: Application to Registration of Abdominal Structures, in: ISBI 2011, IEEE. pp. 1219–1222.

Gottschalk, S., Lin, M.C., Manocha, D., 1996. OBBTree: A Hierarchical Structure for Rapid Interference Detection, in: SIGGRAPH, ACM. pp. 171–180.

Kovan, V., 2008. An Assessment of Impact Strength of the Mandible. Journal of Biomechanics 41, 3488–3491.

Kozic, N., Weber, S., Büchler, P., Lutz, C., Reimers, N., González Ballester, M.A., Reyes, M., 2010. Optimisation of Orthopaedic Implant Design Using Statistical Shape Space Analysis Based on Level Sets. Medical Image Analysis 14, 265–275.

Lovald, S., Baack, B., Gaball, C., Olson, G., Hoard, A., 2010. Biomechanical Optimization of Bone Plates Used in Rigid Fixation of Mandibular Symphysis Fractures. Journal of Oral and Maxillofacial Surgery 68, 1833–1841.

Lovald, S.T., Wagner, J.D., Baack, B., 2009. Biomechanical Optimization of Bone Plates Used in Rigid Fixation of Mandibular Fractures. Journal of Oral and Maxillofacial Surgery 67, 973–985.

Martìn-Fernández, M.A., Martìn-Fernández, M., Alberola-López, C., 2009. A Log-Euclidean Polyaffine Registration for Articulated Structures in Medical Images, in: Yang, G.Z., Hawkes, D.J., Rueckert, D., Noble, J.A., Taylor, C.J. (Eds.), MICCAI 2009, Part I, Springer, Heidelberg. pp. 156–164.

McBroom, R.J., Hayes, W.C., Edwards, W.T., Goldberg, R.P., White, A.A., 1985. Prediction of Vertebral Body Compressive Fracture Using Quantitative Computed Tomography. The Journal of Bone and Joint Surgery 67, 1206–1214.

Metzger, M.C., Vogel, M., Hohlweg-Majert, B., Mast, H., Fan, X., Rüdell, A., Schlager, S., 2011. Anatomical Shape Analysis of the Mandible in Caucasian and Chinese for the Production of Preformed Mandible Reconstruction Plates. Journal of Cranio-Maxillo-Facial Surgery 39, 393–400.

Moore, G., Olson, T., Yonkers, A., 1985. A Retrospective of 100 Fractures in 56 Patients. Nebraska Medical Journal 70, 120–123.

Muscolo, D.L., Ayerza, M.A., Aponte-Tinao, L.A., Ranalletta, M., 2005. Use of Distal Femoral Osteoarticular Allografts in Limb Salvage Surgery. J Bone Joint Surg Am 87, 2449–2455.

Pizer, S.M., Fletcher, P.T., Joshi, S., Thall, A., Chen, J.Z., Fridman, Y., Fritsch, D.S., Gash, A.G., Glotzer, J.M., Jiroutek, M.R., Lu, C., Muller, K.E., Tracton, G., Yushkevich, P., Chaney, E.L., 2003. Deformable M-Reps for 3D Medical Image Segmentation. International Journal of Computer Vision 55, 85–106.

Ritacco, L.E., Seiler, C., Farfalli, G.L., Nolte, L., Reyes, M., Muscolo, D.L., Tinao, L.A., 2012. Validity of an Automatic Measure Protocol in Distal Femur for Allograft Selection from a Three-Dimensional Virtual Bone Bank System. Cell and Tissue Banking , 1–8.

Seiler, C., Pennec, X., Reyes, M., 2011a. Geometry-Aware Multiscale Image Registration Via OBBTree-Based Polyaffine Log-Demons, in: Peters, T., Fichtinger, G., Martel, A. (Eds.), MICCAI 2011, Springer, Heidelberg.

Seiler, C., Pennec, X., Ritacco, L., Reyes, M., 2011b. Femur Specific Polyaffine Model to Regularize the Log-Domain Demons Registration, in: SPIE Medical Imaging.

Siddiqi, K., Pizer, S.M., 2008. Medial Representations: Mathematics, Algorithms and Applications. Springer Verlag.

Taquet, M., Macq, B., Warfield, S.K., 2011. Spatially Adaptive Log-Euclidean Polyaffine Registration Based on Sparse Matches, in: Peters, T., Fichtinger, G., Martel, A. (Eds.), MICCAI 2011, Springer, Heidelberg.

Urken, M.L., Weinberg, H., Vickery, C., Buchbinder, D., Lawson, W., Biller, H.F., 1991. Oromandibular Reconstruction Using Microvascular Composite Free Flaps. Report of 71 Cases and a New Classification Scheme for Bony, Soft-Tissue, and Beurologic Defects. Archives of Otolaryngology–Head & Neck Surgery 117, 733–744.

Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., 2009. Diffeomorphic Demons: Efficient Non-Parametric Image Registration. NeuroImage 45, S61–S72.

Watanabe, H., Abdul, M.M., Kurabayashi, T., Aoki, H., 2010. Mandible Size and Morphology Determined with CT on a Premise of Dental Implant Operation. Surgical and Radiologic Anatomy 32, 343–349.

Zhang, W., Brady, M.J., Becher, H., Noble, A.J., 2011. Spatio-Temporal (2D+T) Non-Rigid Registration of Real-Time 3D Echocardiography and Cardiovascular MR Image Sequences. Physics in Medicine and Biology 56, 1341+.

Zhuang, X., Rhode, K.S., Razavi, R.S., Hawkes, D.J., Ourselin, S., 2010. A Registration-Based Propagation Framework for Automatic Whole Heart Segmentation of Cardiac MRI. IEEE Transactions on Medical Imaging 29, 1612–1625.