

ADR - Anatomy-Driven Reformation

Jan Kretschmer, Grzegorz Soza, Christian Tietjen, Michael Suehling, Bernhard Preim, and Marc Stamminger

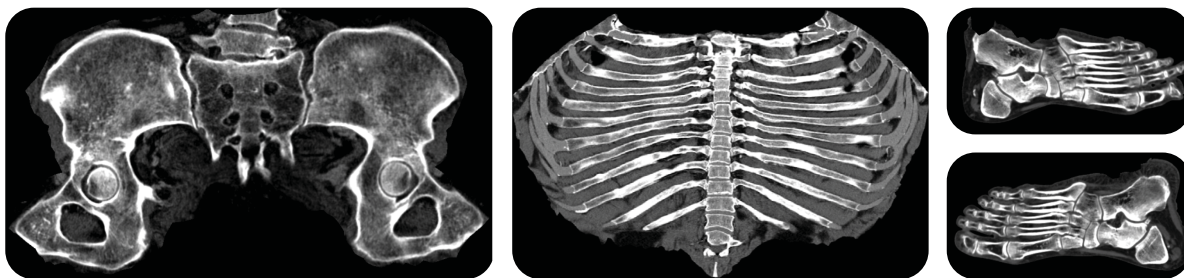


Fig. 1. Several results generated with our Anatomy-Driven Reformation approach (ADR). By taking the individual shape of an anatomical structure into account, our method allows generating flat reformations that provide comprehensive overviews. *From left to right: pelvis ADR, rib cage ADR and feet ADR.*

Abstract—Dedicated visualization methods are among the most important tools of modern computer-aided medical applications. Reformation methods such as Multiplanar Reformation or Curved Planar Reformation have evolved as useful tools that facilitate diagnostic and therapeutic work. In this paper, we present a novel approach that can be seen as a generalization of Multiplanar Reformation to curved surfaces. The main concept is to generate reformatted medical volumes driven by the individual anatomical geometry of a specific patient. This process generates flat views of anatomical structures that facilitate many tasks such as diagnosis, navigation and annotation. Our reformation framework is based on a non-linear as-rigid-as-possible volumetric deformation scheme that uses generic triangular surface meshes as input. To manage inevitable distortions during reformation, we introduce importance maps which allow controlling the error distribution and improving the overall visual quality in areas of elevated interest. Our method seamlessly integrates with well-established concepts such as the slice-based inspection of medical datasets and we believe it can improve the overall efficiency of many medical workflows. To demonstrate this, we additionally present an integrated visualization system and discuss several use cases that substantiate its benefits.

Index Terms—Medical Visualization, Volume Reformation, Viewing Algorithms

1 INTRODUCTION

Computer-generated visualizations play an important role in modern clinical practice since they enable a very flexible and efficient inspection of radiological data. The importance and the positive impact of computer-aided workflows on the overall efficiency of today's radiological practice is well documented by extensive studies [31, 32]. Moreover, some studies already predict an alarming shortage in diagnostic radiologists over a 30 year horizon [3]. Thus, one key objective of medical visualization is to improve the efficiency of radiological work in an increasingly demanding environment. This means that tasks have to be facilitated for faster processing without sacrificing any accuracy or completeness.

Not only in the oncological and the trauma-related field, reconstructions from Computed Tomography (CT) are often the basis for both initial diagnosis and ongoing treatment. Down to the present day, the gold standard in diagnostic radiological practice is a slice-wise inspection of standard orthogonal views. This direct inspection of reconstructed slices is not ideal for many tasks since anatomical structures are generally *not* aligned with the CT coordinate system and may exhibit complex shapes. The detection of rib metastases or fractures, for instance, usually involves the inspection of hundreds of slices in a CT scan. Since standard orthogonal reformations (i.e. axial, sagittal, coronal) provide oblique cuts through ribs, it is a tedious and error prone task to reliably track longitudinal changes as needed for diagnosis [7]. Once detected, metastases have to be documented. This means they have to be tracked back to the corresponding vertebra to allow for a proper anatomical labeling. Using slice-based navigation, this becomes a tedious task since the cross-sections of a rib bone usually expose a high drift between slices. Recent segmentation and visualization methods have started to pick up on this by proposing solutions that provide centerline-driven normalized views to radiologists [43].

Similarly, pelvic bone lesions can exhibit a large variety of appearances making their detection and classification a tedious task [13]. The complex anatomy of the pelvic bone further complicates the assessment of lesion geometry (i.e. size and shape) and anatomical location when using standard orthogonal reformations of the CT data set.

The detection and classification of fractures of the skull and the skull base is an important task when dealing with traumatic head injuries. A user study showed that anatomy-aligned curved projections can increase the fracture detection rate compared to traversal views [34]. In addition, the authors show that detection can be performed about four times faster compared to standard views.

In this paper, we present a generic reformation framework that al-

- Jan Kretschmer is with the Department of Computer Graphics, FAU Erlangen, and Siemens Healthcare Computed Tomography, Forchheim, Germany. E-mail: jan.kretschmer@cs.fau.de.
- Grzegorz Soza is with Siemens Healthcare Computed Tomography, Forchheim, Germany. E-mail: grzegorz.soza@siemens.com.
- Christian Tietjen is with Siemens Healthcare Computed Tomography, Forchheim, Germany. E-mail: christian.tietjen@siemens.com.
- Michael Suehling is with Siemens Healthcare Computed Tomography, Forchheim, Germany. E-mail: Michael.Suehling@siemens.com.
- Bernhard Preim is with the Department of Simulation and Graphics, Otto-von-Guericke University of Magdeburg, Germany. E-mail: bernhard@ovgu.de.
- Marc Stamminger is with the Department of Computer Graphics, FAU Erlangen, Germany. E-mail: marc.stamminger@cs.fau.de.

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014. Date of publication 11 Aug. 2014; date of current version 9 Nov. 2014.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

Digital Object Identifier 10.1109/TVCG.2014.2346405

lows to generate Anatomy-Driven Reformations (ADRs) of complex anatomical structures. The reformatted volumes have a highly normalized character since they map patient-specific curved structures to a consistent flat setting. ADRs do not aim at replacing original reconstructions but rather at facilitating common tasks like navigation or annotation. For instance, linking ADR views with standard orthogonal reconstructions, enables more efficient workflows for common oncological and trauma injury-related tasks without sacrificing any image fidelity. In Section 7 we provide more detailed examples for clinically motivated applications that directly benefit from precomputed ADRs.

2 RELATED WORK

The end-to-end reformation framework, we present in this paper, touches many fields of medical visualization and computer graphics. Hence, in this section we will give a brief overview over the most relevant related work and point to summary articles for further reading.

2.1 Anatomy-Aware Visualization of Medical Data

The generation of meaningful and diagnostically relevant visualizations of medical datasets is a challenging task. Due to the high geometrical complexity of the human body, major problems include mutual occlusions, self-occlusions and general visual clutter. Since the inter-patient variation of anatomical geometry is usually relatively small, an immense variety of visualization algorithms has been developed, many of which explicitly take advantage of the shape of particular anatomical structures. In the following, we give an overview on the field of anatomically motivated projection and reformation techniques. We then discuss some previously proposed end-to-end approaches.

2.1.1 Projection Methods

Projection-based visualization techniques constitute a common way to enhance raw medical data and they have been proposed in many different flavors and for broadly varying tasks. A common approach is to approximate organs or other anatomical structures by geometric primitives like spheres [19], cylinders [28] or even planes [25] that allow for a straight-forward projection of surrounding tissue. These methods are usually tweaked for specific applications and often focus on dedicated transfer function designs. Typical applications include projections of cardiac areas or tumors [33]. With simplicity and intuitiveness being their strength, closed-form primitives often fail to sufficiently capture the geometry of the anatomical structure being projected, leading to distortions and occlusions.

2.1.2 Surface Reformation Methods

To allow for a more flexible inspection of medical datasets, Multiplanar Reformation (MPR) [15] was proposed to resample reconstructed CT volumes in arbitrarily oriented planes. Curved Planar Reformation (CPR) [17] and its derivatives [1, 18, 26] provide even more flexible cuts through datasets that are driven by single geometric centerlines or complex centerline graphs. CPRs are a common tool in vessel visualization since they generate cuts that allow for a thorough inspection of vessel lumen and that include valuable anatomical context. In contrast to MPRs, CPRs are driven by patient-specific anatomical information which ultimately allows condensing more information into a single rotatable 2D view. Saroul et al. present similar reformation results for curved cut surfaces [37]. Their mapping algorithm also allows to preserve distances in fixed directions [36] but is generally limited to cut surfaces. The main problem with current reformation-like approaches is that they usually resort to some kind of projection. Moreover, they usually do not provide an explicit handling for the distortions that arise and only parameterize single cuts instead of volumetric regions.

2.1.3 Visualization Frameworks

Dedicated visualization frameworks have been proposed for several anatomical structures. Colon flattening methods [2, 14, 42] for instance, constitute a prominent example for efforts to improve a tedious and error-prone diagnostic task of modern cancer prevention.

Because of their complex shape and size, colons are extremely challenging to visualize. To preserve the distinct shape of important features like polyps in an expressive manner, colon unfolding methods are usually particularly concerned with parameterization and occlusion issues. To speed up the task of detecting and labeling rib bone lesions, a centerline-driven viewing approach was presented in [43]. The method provides a rotatable CPR for each rib and arranges all CPRs in an atlas-like view. To provide cranial unfoldings, a dedicated method for trauma-related readings was developed in [34]. The algorithm basically drops an elastic grid over the head of a patient to wrap and parameterize the skull bone. Subsequently, the grid is used to compute projections which, as the article confirms in a study, lead to an improved sensitivity for the detection of skull fractures. For the diagnosis of coronary artery diseases, Termeer et al. present the volumetric bulls eye plot [41] as an extension to the well established bulls eye plot method. Their method generates a continuous thick parameterization of the myocardium that allows for different unfoldings.

2.2 Conformal Mappings and Volumetric Deformation

Providing some kind of normalized or flattened view of an originally more articulated structure generally introduces inevitable distortions. There are two related fields, one dealing with mappings between manifolds and one dealing with mappings between volumes. We will refer to these as *surface parameterization* and *volume parameterization* respectively. The general problem of finding low-distortion mappings between objects of different shape is a well-known problem in computer graphics. It appears in texture mapping [22], ambient space warping and surface-based mesh deformation [40] and volumetric mesh deformation [46].

Conformal mappings preserve angles, which is an important property when trying to preserve similarity during mapping. The conformality of a mapping plays an important role in image registration tasks in computer vision [24, 44] and medical image registration [10, 47]. As discussed in [23], however, conformal surface parameterizations still allow for local or global scaling to occur. This leads to a blow-up or a shrinkage of certain regions. To overcome this, the authors of [23] borrow from the As-Rigid-As-Possible (ARAP) deformation paradigm introduced by Sorkine et al. [40] and propose a surface parameterization framework that allows to optimize for a trade-off between local conformity (preservation of angles) and local rigidity (preservation of angles and lengths).

Due to this local rigidity preservation and its efficient nature, the ARAP approach has been adopted for volume-aware interactive deformation methods based on grids [11, 46] and skeletons [45]. Because of the high computational burden, volumetric deformations methods often employ coarse-to-fine proxy grids that use different deformation methods in each level (i.e. higher quality on coarser levels to fast approximations in fine levels) [11]. In [21], volumetric deformation method based on interior radial basis functions (IRBFs) is proposed. The key idea is to use shape-interior distances instead of Euclidean distances which allows deformations to adhere to the interior of a shape. By this, close-by thin structures like limbs can be deformed separately, whereas they would influence each other in classical ambient space deformation based on Euclidean distance metrics.

2.3 Model-Based Medical Segmentation

Medical image segmentation is a very common task that has given rise to a large variety of different approaches. Since the method we propose is based on surfaces that preferably run through the center of anatomical structures, from our perspective segmentation is a pre-processing step. Model-based image segmentation is particularly well suited for our application since usually the approximate shape of an anatomical structure and its medial surface is known beforehand. The basic strategy of model-based methods such as Active Shape Models [6] or Active Appearance Models [5] is to exploit this geometric and intensity-based knowledge to achieve segmentations with inherent plausibility. Usually, parametric models are based on observations drawn from a statistical data base [4, 38]. For segmentation, models are often initialized with the statistical mean of the parameter set.

This parameter set is then optimized until similarity metrics between the appearance of the model and the target dataset reach a minimum. Annotating the statistical model with a center surface, for example using vertex skinning or similar techniques, allows for the simultaneous generation of a center surface fit during segmentation.

Medical image registration [24, 44, 47] produces deformation fields that enable warping between a reference dataset and a target dataset. Using this deformation field, annotations in the reference dataset, like center surfaces can be mapped to the patient-specific setting. Medial surfaces [8, 20] can also serve as a starting point to generate meshes that approximately run through the center of a structure.

3 CONTRIBUTION AND OVERVIEW

Current anatomy-aware reformation and visualization methods tend to focus on one specific anatomical structure [1, 18, 34, 43] and thus lack general applicability. While some methods do enable slicing by moving curved surfaces [36, 37] they do not provide a global minimization of distortion artifacts between slices. Volumetric deformation methods, on the other hand, are usually designed for generality and provide a powerful basis to create low-distortion mappings between structures. Current applications to medical visualization [11], however, lack adaptation to the clinical scenarios that motivated our work, i.e. the creation of anatomy-aligned thick reformations (see Section 7.1).

We propose an extension of the ARAP surface parameterization scheme proposed in [23] to a thick volumetric region around a surface mesh. It allows users to reformat (i.e. *flatten*) the surrounding of cut surfaces through medical datasets and can be viewed as an extension of standard Multiplanar Reformation to curved surfaces. Our method describes a hybrid approach between a single cut surfaces [23] and volumetric deformation [11, 46]. Similar to conventional MPRs our method creates reformations based on a surface. However, since we allow surfaces to be curved, distortions in the deformed region need to be managed during flattening. This is where we draw from the field volumetric deformation to globally minimize distortions.

As a key contribution, our method generates reformations that are aligned with the patient-specific geometry of an anatomical structure. This allows us to create Anatomy-Driven Reformations (ADR) tailored to a large variety of entirely different anatomical structures. Our formulation thus provides a means to integrate many different applications into a single pipeline. We demonstrate this by proposing a reformation-based application that enables a fast exploration of complex bone structures. In addition, our formulation includes an importance-based optimization scheme to manage distortions and to increase the overall visual quality.

The basic input for our method is an *ADR surface*, represented as a triangulated mesh in world space that runs through an anatomical structure of interest. The ADR surface should usually be located in the center of the structure since its surrounding to both sides will be parameterized by the reformation pipeline. The output of our method consists of flat reformatted medical volumes that can be inspected with well known approaches like slicing. To allow for this kind of unfolded anatomy-aligned slicing, the ADR surface needs to be open.

Figure 2 illustrates our basic pipeline and the most important structures involved. Our framework consists of three main steps:

1. **Offset surface computation:** In the first step we compute copies of the surface mesh with negative and positive displacements. The region between the resulting three layers defines the thickness and the content of the final reformation.
2. **Deformation:** We then compute a flat embedding of the three surfaces in three flat stacked layers. The mapping that corresponds to a volume parameterization explicitly minimizes intra-layer distortions and shearing errors between layers.
3. **Reformation:** After the ADR layers are optimized, we use the stacked setting to resample a reformatted volume. The correspondence to world space can be used for atlas-like navigation and annotation.

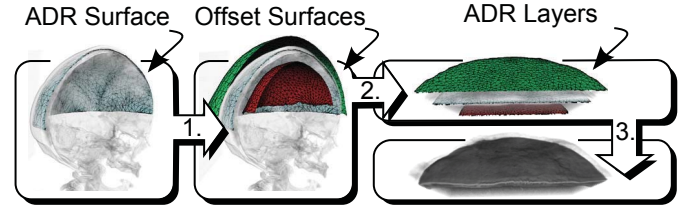


Fig. 2. Basic outline and terminology of our reformation pipeline using the example of a skull reformation (renderings are clipped for clearer visibility). Offset surfaces need to be computed for the initial ADR surface to parameterize its surrounding volume (1). By deforming the resulting mesh to a flat layered setting, a mapping is induced (2). Resampling the original dataset results in a reformatted flat volume (3).

4 ANATOMY-DRIVEN REFORMATION (ADR)

The key concept of our visualization method is to generate reformatted views of medical datasets based on a given 2D ADR surface embedded in 3D space. In this section we will first lay out the pipeline for a surface-based reformation framework which corresponds to the parameterization method described in [23]. We will then extend the method to a volumetric approach that generates a parameterization for the surrounding area of the ADR surface.

4.1 Input Data

The ADR surface mesh \mathbf{M} is described by a list of vertices \mathbf{V} with coordinates $\mathbf{v}_i \in \mathbb{R}^3, i = 0 \dots n-1$, a list of triangles \mathbf{T} and a list of half edges \mathbf{HE} . To be able to generate flattened versions of 3D meshes, one restriction we impose is that the meshes are not closed, i.e. that they have at least one border so they can be flattened. In practice, this restriction does not pose significant problems, since closed meshes can be cut open interactively or automatically beforehand [39].

4.2 Surface-Based ADR

When reformatting a medical volume, the goal is to map \mathbf{M} to a flat representation \mathbf{M}' with vertices $\mathbf{v}'_i \in \mathbb{R}^2$. This introduces distortions which should be minimized to provide a reformation that preserves the anatomical geometry as much as possible. For this, we extend the mesh parameterization method described in [23]. In this approach, the surface mesh is viewed as a set of connected cells. Each cell corresponds to a triangle of the mesh. This means each cell can be embedded in a plane trivially since triangles by definition only span a 2D subspace. When computing the mapping from the mesh embedded in \mathbb{R}^3 to the planar embedding in \mathbb{R}^2 , the key idea is to keep every individual cell as rigid as possible while at the same time forcing the shared vertices of the cells (i.e. the triangle corners) to coincide.

To achieve this, we start off with an initial flat embedding \mathbf{M}'_0 of \mathbf{M} and use an iterative non-linear optimization scheme that uses two phases per iteration, as proposed in [23]. In the first (*local*) optimization phase, for every triangle a 2D rotation R_t needs to be computed that best rotates the original undeformed triangle to the current solution. The second (*global*) phase takes care of stitching the individually rotated triangles t_i to a mesh with globally consistent vertex positions. This is necessary since the optimal rotations are computed per triangle and are thus not consistent across the mesh.

4.2.1 Initial Flat Triangulation

The iterative optimization scheme needs an initial flat embedding of the mesh in 2D. As stated in [23] there are several interchangeable ways to compute this initial solution for the vertex positions \mathbf{v}'_i . Because of its simplicity we use an approach that computes an initial harmonic triangulation, similar to the one proposed by Floater [12]. For this we collect all boundary vertices of the mesh and distribute them in consecutive order along the circumcircle of a disc. We then use these outer vertices as boundary conditions and solve the following

discretized linear Poisson system for all free vertices $\mathbf{v}'_i, i \notin \Omega$.

$$\#V(i)\mathbf{v}'_i - \sum_{j \in V(i) \cap \bar{\Omega}} \mathbf{v}'_j = \sum_{j \in V(i) \cap \Omega} \mathbf{v}'_j, \forall i \notin \Omega \quad (1)$$

Here, $V(i)$ denotes the set of direct neighbors of a vertex \mathbf{v}_i and $\#V(i)$ is the size of the set (i.e. the valence of the vertex). Ω is the set of constrained boundary vertices which are on the right-hand side of Equation 1 and which are known. Because of the Laplacian mean value property, the resulting vertex positions lead to a flat mesh that does not fold and thus constitutes a well-behaved initial solution for the subsequent iterative optimization scheme. The initialization process is illustrated in Figure 3.

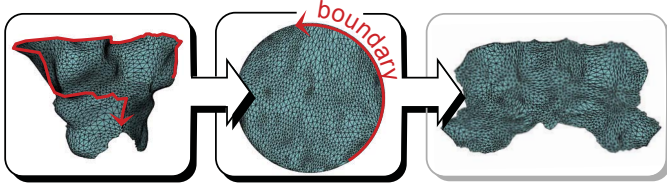


Fig. 3. *Left*: 3D ADR Surface for a pelvis reformation with some boundary vertices marked. *Center*: Initial flat triangulation. Boundary vertices are distributed onto a circle and serve as boundary conditions for harmonic vertex positions based on the connectivity of the mesh. *Right*: Final flat triangulation after 20 iterations.

4.2.2 Local Phase - Computing Triangle Rotations

Following [23], the purpose of the local phase is to explicitly compute a 2D similarity transformation or a rigid transformation $R_t: \mathbb{R}^2 \rightarrow \mathbb{R}^2$. The purpose of R_t is to align the original, undeformed triangle to the geometry of the current solution. To prevent scaling-induced distortions from creeping into our final mapping, we only allow rotations. This means R_t constitutes an isometric mapping that preserves the area and the angles of each original triangle. To compute rotations, we first need to find a shape-preserving 2D embedding of all original triangles. For this we create an orthonormal basis matrix aligned with each triangle. We then transpose (i.e. invert) it and drop the superfluous third component to get a mapping to the triangle tangent space $N_t: \mathbb{R}^3 \rightarrow \mathbb{R}^2$

$$N_t(\mathbf{x}) = \left(\frac{\mathbf{v}_b - \mathbf{v}_a}{\|\mathbf{v}_b - \mathbf{v}_a\|} \frac{\mathbf{n} \times (\mathbf{v}_b - \mathbf{v}_a)}{\|\mathbf{n} \times (\mathbf{v}_b - \mathbf{v}_a)\|} \right)^T \left(\mathbf{x} - \frac{1}{3}(\mathbf{v}_a + \mathbf{v}_b + \mathbf{v}_c) \right), \quad (2)$$

$$\mathbf{n} = \frac{(\mathbf{v}_c - \mathbf{v}_a) \times (\mathbf{v}_b - \mathbf{v}_a)}{\|(\mathbf{v}_c - \mathbf{v}_a) \times (\mathbf{v}_b - \mathbf{v}_a)\|}$$

where a, b and c denote the vertex indices of triangle \mathbf{t}_i . For every triangle \mathbf{t}_i we can use the mapping N_t to create an arbitrarily rotated 2D embedding \mathbf{t}_i^0 . To find the optimal rotation between this original triangle $(\mathbf{v}_a^0, \mathbf{v}_b^0, \mathbf{v}_c^0)_t = (N_t(\mathbf{v}_a), N_t(\mathbf{v}_b), N_t(\mathbf{v}_c))$ and the deformed triangle in the current solution $(\mathbf{v}'_a, \mathbf{v}'_b, \mathbf{v}'_c)_t$ we use Kabsch's algorithm [16]. For this we need to compute the singular value decomposition $A = VSW^T$ of the covariance matrix A of the centroid-aligned triangle vertices.

$$A = \mathbf{v}_a^0(\mathbf{v}'_a - \mathbf{c})^T + \mathbf{v}_b^0(\mathbf{v}'_b - \mathbf{c})^T + \mathbf{v}_c^0(\mathbf{v}'_c - \mathbf{c})^T$$

where $\mathbf{c} = \frac{1}{3}(\mathbf{v}'_a + \mathbf{v}'_b + \mathbf{v}'_c)$ is the centroid of the current triangle. The rotation is then defined as

$$R_t = W \begin{pmatrix} 1 & 0 \\ 0 & \text{sign}(\det(WV^T)) \end{pmatrix} V^T$$

4.2.3 Global Phase - Solving for Vertex Positions

To find consistent vertex positions for the current set of rotations R , we minimize the following energy functional taken from [23]

$$E_{ARAP}(\mathbf{v}', R) = \frac{1}{2} \sum_{(i,j) \in \mathbf{HE}} \cot(\Theta_{i,j}) \|(\mathbf{v}'_i - \mathbf{v}'_j) - R_{t(i,j)}(\mathbf{v}_i^0 - \mathbf{v}_j^0)\|^2 \quad (3)$$

where $\cot(\Theta_{i,j})$ are cotangent weights as presented in [30], \mathbf{HE} is the set of half edges in the mesh, $R_{t(i,j)}$ is the rotation for the triangle associated with half edge (i, j) , $\mathbf{v}_i^0 \in \mathbb{R}^2$ are the vertex positions of the original triangle embeddings and $\mathbf{v}'_i \in \mathbb{R}^2$ are the target vertex positions. By fixing the rotations to the ones computed in the local phase, E_{ARAP} becomes linear. The global phase thus boils down to taking the average rotation of the adjacent triangles of an edge, rotating the edges of the original embedding in 2D and finding vertices \mathbf{v}'_i whose edges minimize the deviation from these (prescribed) edges in a least squares sense. Setting the gradient of the energy in Equation 3 with respect to \mathbf{v}'_i to zero results in the following set of linear equations

$$\sum_{j \in V(i)} (\cot(\Theta_{i,j}) + \cot(\Theta_{j,i}))(\mathbf{v}'_i - \mathbf{v}'_j) = \sum_{j \in V(i)} (\cot(\Theta_{i,j})R_{t(i,j)} + \cot(\Theta_{j,i})R_{t(j,i)})(\mathbf{v}_i^0 - \mathbf{v}_j^0) \quad (4)$$

$$\forall i = 1 \dots n - 1$$

Repeatedly computing rotations and then solving the corresponding linear system for the vertex positions \mathbf{v}'_i results in a consistent 2D mesh where every triangle is as close to its original shape as possible.

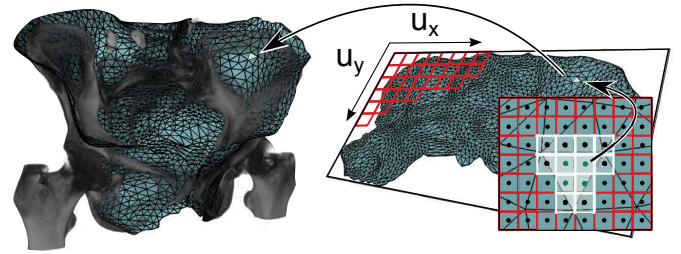


Fig. 4. *Top*: Volume rendering of a pelvic CT scan and the ADR surface mesh in 3D. *Bottom*: The flat parameterization of the surface mesh in ADR image space. For sampling, the marked triangle was enlisted at all highlighted pixels. The green points are the pixels that are actually included in the triangle. They are mapped to the 3D setting to sample draw intensity samples from the original dataset.

4.2.4 Resampling

Once the final 2D vertex positions \mathbf{v}'_i are computed, in conjunction with the original mesh embedded in 3D, they allow for a piece-wise linear mapping between the domain of the dataset and the 2D domain. To generate an ADR image of a certain target resolution $ADR_{resx} \times ADR_{resy}$, we compute axis-aligned bounds \mathbf{v}'_{\min} and \mathbf{v}'_{\max} for the 2D embedding of the mesh and map the coordinates to the ADR image space

$$u_i = \begin{pmatrix} u_{ix} \\ u_{iy} \end{pmatrix} = \begin{pmatrix} (\mathbf{v}'_{ix} - \mathbf{v}'_{\min x}) \frac{ADR_{resx}}{\mathbf{v}'_{\max x} - \mathbf{v}'_{\min x}} \\ (\mathbf{v}'_{iy} - \mathbf{v}'_{\min y}) \frac{ADR_{resy}}{\mathbf{v}'_{\max y} - \mathbf{v}'_{\min y}} \end{pmatrix}. \quad (5)$$

To determine the intensity (or color) value of a pixel in the ADR image, we first need to find the triangle of the flattened mesh that covers the respective pixel. After this we compute the barycentric coordinates of the pixel center to transform it to the 3D domain using the original mesh coordinates \mathbf{v}_i . Sampling the volume at the resulting position gives the intensity value for the corresponding pixel in the ADR image. To speed up the search task of finding the triangle that covers a particular pixel, we compute a lookup table. For this we take the axis-aligned bounds of every triangle and enlist the triangle in a per-pixel list of every pixel covered by the bounds. Using the per-pixel list during sampling, barycentric coordinates only have to be computed for a small local subset of the triangles. Figure 4 illustrates the ADR image sampling process.

4.3 Volumetric Reformation

Creating an ADR image as described in the previous section only generates a reformatted 2D image of a singular cut surface. When inspecting volumetric medical datasets, a common task is to examine the surroundings of a particular site by viewing several adjacent slices of the reconstructed image. This is commonly referred to as *slicing*. To provide a flexible visualization system, that is not limited to a single surface in 3D we extend our method to parameterize not only the ADR surface, but also the surrounding volume.

4.3.1 Normal-Based Volume Parameterization

One way to parameterize the surrounding is to compute offset surfaces. Computing offset surfaces is a common task in computer graphics [29, 35]. The most obvious way to generate offset surfaces is to displace the mesh vertices in the direction of the normal. To generate a volume parameterization for the surrounding of the ADR surface, we compute two displaced copies of the original mesh, one with an offset in negative and one in positive direction of the normal $\mathbf{v}^- = \mathbf{v} - d\mathbf{n}$ and $\mathbf{v}^+ = \mathbf{v} + d\mathbf{n}$. Here, d denotes the offset distance which corresponds to the maximum distance that can be inspected in the final reformation. To generate reformations with consistent thickness, we assume d to be constant across the whole ADR surface which results in a stack of three layered meshes as illustrated in Figure 5.

A straight-forward way to extend the surface-based reformation approach (see Section 4.2) to a volumetric approach, is to flatten the center ADR layer only and to apply the same flat geometry to the offset surfaces. This keeps the reformatted triangle shapes constant in the layer offset dimension and leads to a stack of extruded triangles as illustrated in Figure 5. Analogously to the sampling strategy presented in Section 4.2.4, this setup creates a mapping between the reformatted space and original 3D space which is, however, volumetric.

One problem with this approach is that global and local self-intersections frequently occur when computing the displaced offset meshes. To reduce self-intersections, we apply several iterations of a Laplacian smoothing kernel [27] to the displaced vertex positions. The mean-value property of the Laplacian usually resolves local self-intersections and leads to generally smoother offset meshes. However, the displaced vertices will no longer lie on the normal-rays of the vertices of the ADR surface, which will lead to distortions when using a constant stacked reformation approach as depicted in Figure 5.

4.3.2 As-Rigid-As-Possible Volume Parameterization

Since ADR surfaces are usually not planar, the size of displaced triangles in the 3D setting differs from the size of their original counterparts (see Figure 6, left). The amount of shrinkage or growth depends on the local curvature of the ADR surface and thus varies across the mesh. This means that using the same flat geometry for all ADR layers as described in the previous section will introduce a volumetric distortion and thus violate the local rigidity property (see Section 4.2) of our reformation (see Figure 6). To explicitly account for the size of the triangles of the offset surfaces, we extend the layered mesh \mathbf{M}' to include vertices for all three layers

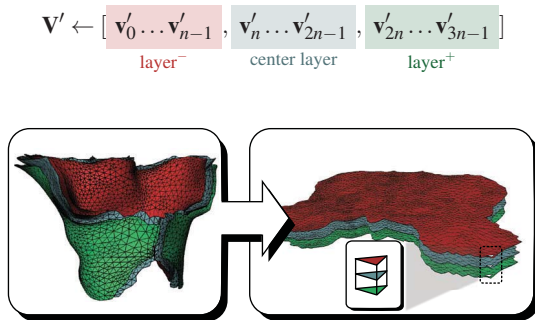


Fig. 5. Left: ADR surface (blue) and the positive (green) and negative (red) offset meshes. Right: The reformatted setting for volumetric reformation with three ADR layers corresponding to the offset meshes.

and extend the topological information (i.e. \mathbf{T}' and \mathbf{HE}') accordingly. This triples the number of unknowns of the resulting linear system (Equation 4) and results in three flat layered triangulations with identical connectivity.

The individual triangulations are flattened with a minimized deviation from rigidity per triangle. This means the positive layer, the center layer and the negative layer provide a distortion-minimizing mapping from ADR image space to the respective surfaces in the 3D volume. Because of the shrinkage and growth of the offset surfaces, however, exclusively minimizing intra-layer rigidity leads to shearing artifacts as illustrated in Figure 6, right.

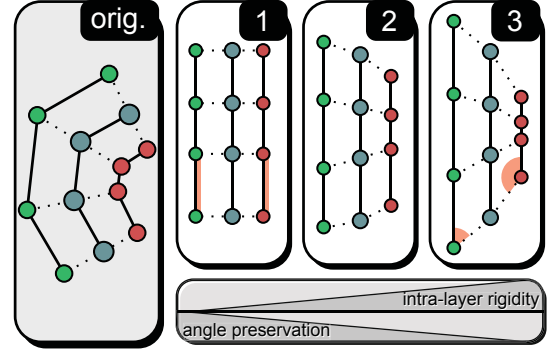


Fig. 6. 2D Illustration of layered volumetric reformation. The original setting is displayed on the left. Using the same flat geometry for all layers leads to good preservation of angles (1) but changes the size of the cells. Explicit ARAP flattening of the individual layers preserves the size of cells, but leads to angular distortion between layers (i.e. shearing) (3). Our method allows for a controllable trade-off between angular and rigidity error in a reformation (2).

To alleviate shearing, we introduce an additional shearing term in the energy functional (recall Equation 3) by adding constraints that relate the individual layers with prescribed vertex offsets. Their purpose is to keep the vertices of the displaced layers and the center layer in similar relative locations with respect to the local tangent space. For this, we use the tangent space transform N_t to transform the offset vertices \mathbf{v}_i^- and \mathbf{v}_i^+ to the tangent space of the ADR surface at the corresponding vertices \mathbf{v}_i . Since N_t is defined for triangles and there are usually several triangles incident to a vertex, we average the offsets contributed by the individual triangles to receive mean vertex offsets \mathbf{o}_i^- and \mathbf{o}_i^+ .

$$\begin{aligned} \mathbf{o}_i^- &= \frac{1}{\#T(i)} \sum_{t \in T(i)} R_t \left(N_t(\mathbf{v}_i^-) - N_t(\mathbf{v}_i) \right) \\ \mathbf{o}_i^+ &= \frac{1}{\#T(i)} \sum_{t \in T(i)} R_t \left(N_t(\mathbf{v}_i^+) - N_t(\mathbf{v}_i) \right) \end{aligned} \quad \forall i = 0 \dots n-1 \quad (6)$$

Here, $T(i)$ is the set of triangles incident at vertex i of the center layer. The highlighted expressions only depend on the geometry of the input ADR surface and are thus constant during optimization.

Introducing the shearing term E_{SH} yields the following energy functional

$$E^\alpha(\mathbf{v}'_i, R) = E_{ARAP}(\mathbf{v}'_i, R) + \alpha E_{SH}(\mathbf{v}'_i, R) \quad (7)$$

where

$$E_{SH}(\mathbf{v}'_i, R) = \sum_{i=0 \dots n-1} \left(\|\mathbf{v}'_{i+n} - \mathbf{v}'_i\| - \mathbf{o}_i^- \|^2 + \|\mathbf{v}'_{i+2n} - \mathbf{v}'_i\| - \mathbf{o}_i^+ \|^2 \right). \quad (8)$$

Here, the weight $\alpha > 0$ allows increasing the influence of the shearing term. Equation 8 essentially constrains the vertices of the displaced layers to expose offsets to their center vertex which are similar to the ones in the undeformed setting with respect to the local tangent space. Thus, for a large α , the reformatted volume will expose less shearing

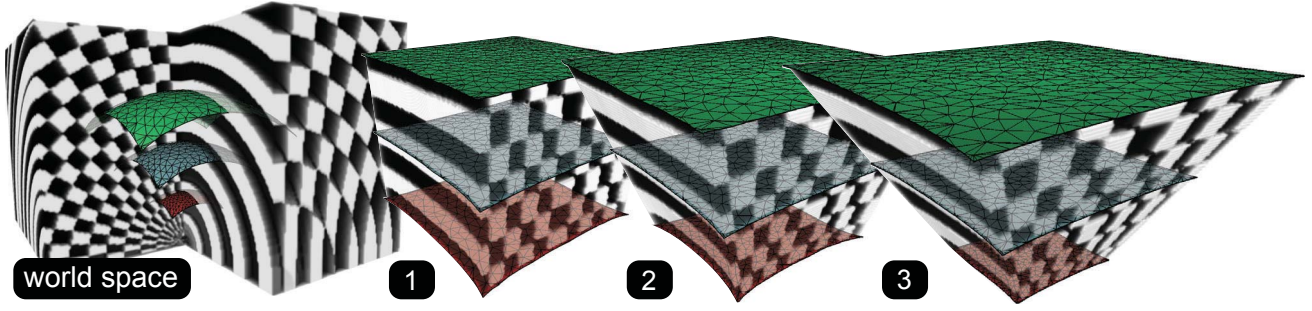


Fig. 7. Reformation results for different weights α for the shearing term E_{SH} . The left shows the original synthetic volume with a spherical checkerboard-like pattern, an ADR surface (blue) and the offset surfaces (green and red). The ADR surface was generated by clipping the surface of a sphere with a cube. Reformation with $0 < \alpha$ results in a good preservation of angles, but the volume of the upper/lower regions is decreased/increased which means rigidity is not preserved (1). This is equivalent to parameterizing the surrounding of the ADR surface along the normals. Removing the weight of the shearing energy (i.e. $0 \lesssim \alpha$) results in a high intra-layer-rigidity but increases shearing (3). Intermediate values for α allow for a reasonable trade-off between intra-layer-rigidity and angular preservation. Setting $\alpha = 0.1$ usually leads to reasonable results (2).

but intra-layer-rigidity decreases. Figure 7 shows the effect of different values for α when reformatting a synthetic volume with a clean spherical checker board parameterization.

Since we assume the layer offset in world space to be constant, as described in Section 4.3.1, E^α does not account for the distance between the reformatted layers, but operates on 2D offsets. This means the vertex ranges $\mathbf{v}'_{i=0\dots n-1}$ and $\mathbf{v}'_{i=2n\dots 3n-1}$, that correspond to the positive and negative layer of the solution, need to be displaced after solving by $(0, 0, d)$ and $(0, 0, -d)$ to span a volumetric slab.

Since the geometry now varies for the different layers, a slight modification to the sampling strategy (see Section 4.2.4) and the mapping between ADR space and world space is required. Before computing barycentric coordinates for the current triangle, it needs to be interpolated since it is now dependent on the z position in the ADR image. This interpolation can be carried out with linear blending. The triangle at a certain height in the layered ADR stack is thus defined as follows

$$T(z) = u(\mathbf{v}'_{a-n}, \mathbf{v}'_{b-n}, \mathbf{v}'_{c-n}) + v(\mathbf{v}'_a, \mathbf{v}'_b, \mathbf{v}'_c) + w(\mathbf{v}'_{a+n}, \mathbf{v}'_{b+n}, \mathbf{v}'_{c+n}), \\ z \in [0, 2d]$$

$$(u, v, w) = \begin{cases} (1 - \frac{z}{d}, \frac{z}{d}, 0) & 0 < z < d \\ (0, 1 - \frac{z-d}{d}, \frac{z-d}{d}) & d < z < 2d \end{cases} \quad (9)$$

where $(\mathbf{v}'_{a-n}, \mathbf{v}'_{b-n}, \mathbf{v}'_{c-n})$, $(\mathbf{v}'_a, \mathbf{v}'_b, \mathbf{v}'_c)$ and $(\mathbf{v}'_{a+n}, \mathbf{v}'_{b+n}, \mathbf{v}'_{c+n})$ are the instances of a triangle in the negative, center and positive ADR layers respectively. Since triangles may now span different pixels depending on z , we merge the axis-aligned bounding boxes of a triangle in all layers when enlisting them for sampling as described in Section 4.2.4.

5 IMPORTANCE MAPS FOR ANATOMY-DRIVEN REFORMATION

As laid out in Sections 4.2 and 4.3 the reformation process (surface-based or volumetric) introduces inevitable distortions if the ADR surface is not coincidentally developable. Since usually certain parts of the ADR surface run through areas of higher interest than others, it is desirable to account for this during reformation. For this, we introduce *importance maps* to control the distribution of the error during reformation. The idea is to reduce the error in areas of high importance by admitting higher distortions in areas of low importance.

To represent importance maps, we allow the vertices of the ADR layers to be annotated with an additional weight w_i that corresponds to the importance of the area surrounding to the vertex. This can be done manually by annotating an ADR template mesh beforehand, or interactively by reassigning (i.e. painting) weights in the reformatted view and restarting the reformation process. In applications like bone reformations based on CT images, approximate bone masks can easily be generated automatically by thresholding and post-processing the input image. These masks can directly be used to assign importance weights to the vertices of an ADR surface to distortions at bone regions.

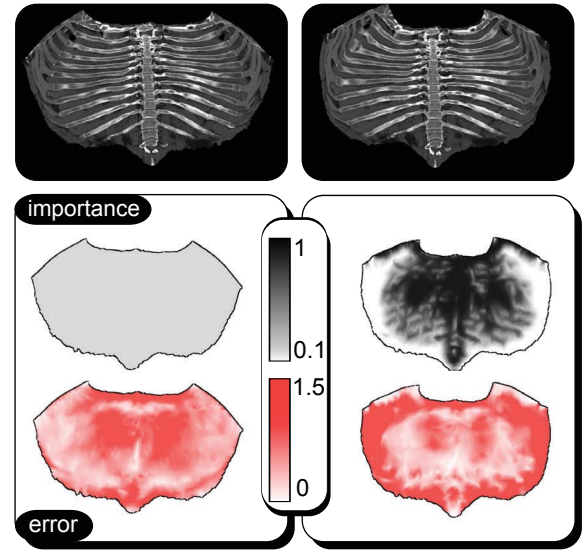


Fig. 8. Two reformations based on the same ADR surface with homogeneous importance (left) and an importance map that emphasizes ribs (right). The importance map in the right reformation allows a higher distortion in the area of the sternum and the cartilage connecting it to the ribs. This results in an overall reduction of the unweighted error in area of the ribs (bottom row). The average weighted error is 7.40% in the constant importance reformation and 6.49% in the importance-driven reformation which corresponds to an error reduction of $\approx 12\%$. The uniform error of the high-importance area drops from 7.90% to 5.69% which corresponds to an error reduction of $\approx 28\%$ in that area (see Table 1).

To incorporate importance maps, we weight every term in the sum of Equation 3 by the importance of the respective half edge $w_{(i,j)}$ which we take as the average of its end vertex weights $w_{(i,j)} = 0.5 \cdot (w_i + w_j)$. The final system of equations for the global phase that includes shearing and importance energy is thus defined as

$$\sum_{j \in V(i)} w_{(i,j)} (\cot(\Theta_{i,j}) + \cot(\Theta_{j,i})) (\mathbf{v}'_i - \mathbf{v}'_j) + A_i \\ = \sum_{j \in V(i)} w_{(i,j)} (\cot(\Theta_{i,j}) R_{t(i,j)} + \cot(\Theta_{j,i}) R_{t(j,i)}) (\mathbf{v}^0_i - \mathbf{v}^0_j) + B_i \quad (10) \\ \forall i = 0 \dots 3n-1$$

Here A_i and B_i originate from the shearing term and connect the layers.

$$(A_i, B_i) = \begin{cases} (\mathbf{v}'_i - \mathbf{v}'_{n+i}, -\mathbf{o}_i^-) & i < n \\ (2\mathbf{v}'_i - \mathbf{v}'_{i-n} - \mathbf{v}'_{i+n}, \mathbf{o}_{i-n}^- + \mathbf{o}_{i+n}^+) & n < i < 2n \\ (\mathbf{v}'_i - \mathbf{v}'_{i-n}, -\mathbf{o}_{i-2n}^+) & 2n < i < 3n \end{cases} \quad (11)$$

The left hand side of Equation 10 is static and the corresponding linear system is sparse and symmetric. This means it can be factored once per reformation and can be reused during iterations. The right hand side has to be recomputed after each iteration since the rotations R_i depend on the geometry of the current solution.

6 RESULTS

The versatility of our approach allows for the creation of reformations for a large variety of anatomical structures. Figure 9 shows several anatomical structures with the corresponding ADR surface and the final reformation results. It features reformations for feet, a pelvis, a skull and a rib cage. To place the ADR surfaces into anatomical structures, we initialize and fit a morphable statistical model similar to Seim et al. [38]. We defined the ADR surface in the coordinate system of the mean shape of the statistical model and deformed the surface during model fitting using standard linear vertex skinning.

The use of the as-rigid-as-possible deformation framework results in low distortions (i.e. deviations from local rigidity) for reformatting volumes. To quantify the error of a reformation, we take the relative errors $\text{Error}_{(i,j)}$ of all individual edges and compute the average over the whole mesh

$$\begin{aligned} \text{Error}_{(i,j)} &= \frac{\text{abs}(\|\mathbf{v}'_i - \mathbf{v}'_j\| - \|\mathbf{v}_i - \mathbf{v}_j\|)}{\|\mathbf{v}_i - \mathbf{v}_j\|} \\ \text{Error} &= \frac{\sum_{(i,j) \in \text{HE}} w_{(i,j)} \text{Error}_{(i,j)}}{\sum_{(i,j) \in \text{HE}} w_{(i,j)}}. \end{aligned} \quad (12)$$

By weighting edge errors according to the importance weights of the corresponding edge $w_{(i,j)}$ as proposed in Equation 12, we get a total error that penalizes distortions in important areas and allows for higher distortions in unimportant areas. Setting $w_{(i,j)} = 1$ when computing the error gives an unbiased error. Figure 8 shows that by supplying importance maps, the average error of a reformation can be significantly reduced. In the example, an error reduction of $\approx 12\%$ is achieved and the errors in the spinal area and the rib area are visibly pushed to the boundary regions of the reformation. An error of 10% means that when measuring lengths of 4 cm in the reformatting volume, the actual length in the curved world space setting is 4 mm off on average. Since we chose the reformation thickness to enclose the bone structures, the offset layers are mostly situated in surrounding soft tissue. For this, we set the importance weights based on the center layer only and copied them to the corresponding vertices in the offset layers, since otherwise these layers would receive low importance and higher error. This is however due to our application and not a general restriction.

Noticeable errors occur when distortions simply cannot be avoided entirely during flattening. This behavior is visible in the vicinity of the upper spine in Figure 8 (right, center). Table 1 lists error totals for the examples presented in Figure 9 and timings for the whole reformation process. It shows that in general the error is low and that importance maps help to significantly reduce distortions during reformation. We list the errors for constant weights and for automatically derived bone importance weights (see Section 5) averaged over the whole mesh. Here, bone regions were assigned importance weight 1.0 and non bone regions were assigned 0.1. In addition, we computed separate non-weighted error averages over the none bone regions (−) and the bone regions (+). Table 1 shows how importance maps allow to distribute the error from important regions to unimportant ones. For all examples except the skull, the total weighted error was reduced by importance maps this is because the entire ADR surface is situated inside of the bone and thus all edges receive the same weight. For importance driven reformations, the uniform error measure, drops in (+) at the expense of higher errors in (−). Since our ADR surfaces are in

Table 1. Errors and timings for the reformations shown in Figure 9. Based on bone masks, all ADR surfaces were split in none bone (−) and bone regions (+). The average error for the whole mesh was computed on reformations with constant importance maps ($E_{\text{const}}^{\text{total}}$) and image-based importance maps ($E_{\text{weighted}}^{\text{total}}$). In the latter case, the importance weight was set to 0.1 for none bone and to 1.0 for bone regions. For both settings (constant importance and varying importance) we additionally computed errors with constant weights for both regions (i.e. E_{const}^- and E_{const}^+). The error of the skull reformation does not decrease since the entire ADR surface is inside the skull bone (− is empty). The high error for the skull reformation is due to its sphere-like geometry resulting in higher distortions for flat embeddings. 100 ARAP iterations were performed for each dataset. The total reformation time T^{total} and the times for optimization T^{opt} and resampling T^{saml} are listed separately (taken on an Intel Core i7 with 4 cores @2.30 GHz and 8GB RAM).

Dataset (# Vertices)	$E_{\text{const}}^{\text{total}}$ (%)		$E_{\text{weighted}}^{\text{total}}$ (%)		T^{total} (s)	
	E_{const}^-	E_{const}^+	E_{const}^-	E_{const}^+	T^{opt}	T^{saml}
Feet (1489)	3.19		1.55		7.50	
	3.62	2.51	4.68	1.01	4.03	3.467
Pelvis (1977)	3.93		3.16		14.54	
	4.21	3.76	6.40	2.89	5.85	8.69
Head (1182)	11.07		11.07		15.07	
	0	11.07	0	11.07	4.72	10.35
Rib Cage (1568)	7.40		6.49		11.76	
	6.86	7.90	11.40	5.69	4.10	7.66

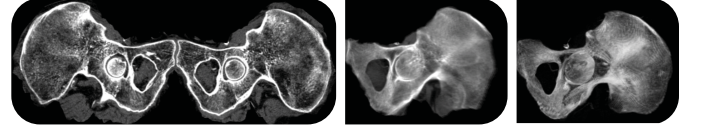


Fig. 10. Different projections of an alternative ADR for a pelvic bone (compare to Figure 9 that includes the sacrum). *Left*: no projection. *Center*: mean intensity projection. *Right*: maximum intensity projection. Intensity windowing functions were adjusted for consistent brightness.

general well adjusted to the anatomical geometry, the bone regions are usually larger than the non bone regions. This explains, why the total errors are more similar to the error of the bone regions than to the error of the non bone regions.

For most anatomical structures, multiple possibilities exist for the shape of a reasonable ADR surface. When reformatting a pelvic bone for instance, the sacrum might not be of interest or when examining a fracture, a reformation of the left or the right side of the pelvis might be sufficient. Figure 10 shows an alternative reformation of a pelvis with an increased focus on the ilium and the pubis and that detaches the pelvis at the sacrum rather than at the pubis (compare Figure 9).

7 APPLICATIONS FOR ANATOMY-DRIVEN REFORMATINGS

For an efficient use of Anatomy-Driven Reformations, it is crucial to embed the reformatting volume into an adequate medical visualization environment. In this section, we will first discuss several use-cases that in large parts motivated the development of the presented method. Based on these applications, we will then lay out a proposal for a reformation platform designed to speed up the corresponding diagnostic and navigational tasks.

7.1 Clinical Use-Cases

Upon the detection of malignant primary tumors, it is usually mandatory for radiologists to examine the skeleton of a patient for secondary metastases [7]. This is a time-consuming task that has to be performed in regular intervals. Similarly, the presence of trauma injuries requires a thorough inspection of the respective affected structure.

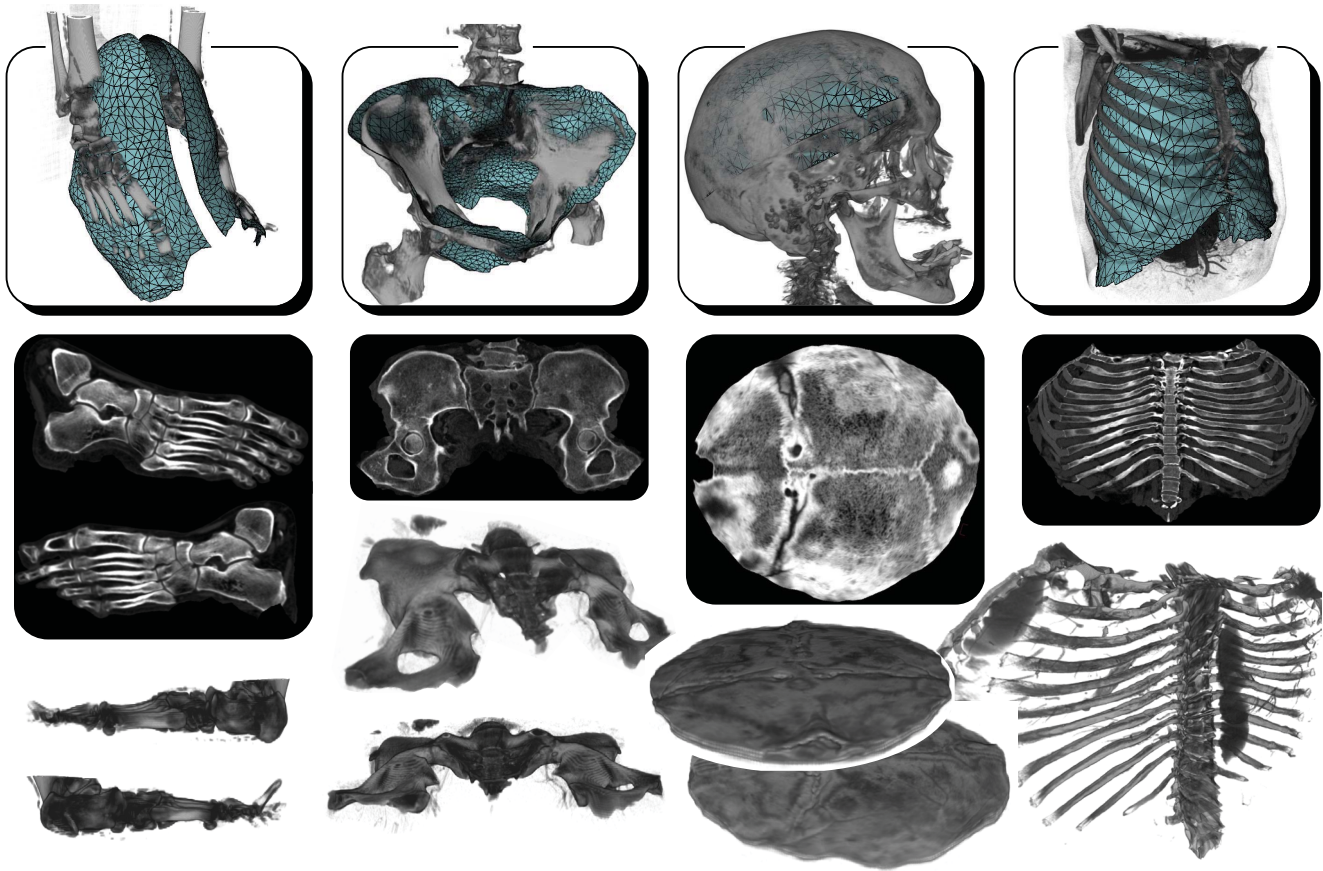


Fig. 9. Several examples for volumetric anatomy-driven bone reformations, where each column is associated with one anatomical structure. The top image of each column shows a volume rendering of the anatomical structure in its original shape and context. The center shows a slice from the reformatted volume. The bottom shows several renderings of the reformatted volumes that illustrate the normalizing character of our method.

Rib Bone Lesions In case of pathologies in the rib cage, clinical documentation usually requires findings to be assigned to the corresponding anatomical rib number (T1-T12). When working in standard orthogonal views, this can be a tedious and error prone task, since any detection in a rib needs to be tracked back to the corresponding vertebra. After this, the spine needs to be examined to determine the vertebra number. The same task has to be performed in the presence of rib fractures.

Pelvic Bone Lesions As in the rib bone scenario, the assessment of metastases in the pelvic bone is generally aggravated by its curved geometry. Depending on the type and progression of a tumor, it can be enclosed entirely by the bone making its geometry depending on the curvature of the Pelvis.

Skull fractures The detection and assessment of skull fractures is an important task in trauma imaging since they can lead to severe complications like cerebrospinal fluid leakage or delayed complications even if no instant hematoma is present [34]. Here, the goal is to detect fractures and to scan their surroundings for abnormalities. Planar cuts do not allow to capture larger regions of the skull surface reducing the visibility of clear large-scale crack patterns. This complicates the detection of thin fractures which is already inherently error prone [34]. Moreover, the curved shape of the skull complicates length measurements which are helpful in fracture classification.

For the tasks listed above, ADRs are particularly well suited since they enable an inspection of a whole anatomical structure in significantly fewer 2D images as compared to standard slicing. This is because the thickness of an ADR (and thus the number of slices to be investigated) is determined by the thickness of the structure under inspection and not by the extents of its axis-aligned bounds in voxel coordinates. In case of the rib cage reformation shown in Figure 9, for

instance, the structure covers $\approx 350\text{mm}$ in height which corresponds to 700 axial slices for the present voxel spacing of $\Delta_x = \Delta_y = \Delta_z = 0.5\text{mm}$. The ADR volume we present, in contrast, has a thickness of $\approx 100\text{mm}$. When using the same voxel spacing as in the original image the reformation has 200 slices. To avoid sampling artifacts we usually slightly oversample the original image. Even though the individual slices contain more voxels (their resolution can be derived from the voxels covered by the ADR surface in world space) the overall amount of image data to be inspected and thus the time needed for examination are reduced since unimportant areas of the volume are not reformatting.

The condensed anatomical information of ADR images additionally allows users to exploit symmetries that are common in many anatomical structures. This helps to increase the sensitivity for anomalies that might be pathological (see ribs and pelvis in Figure 11). Moreover, ADRs facilitate annotation tasks like the determination of the anatomical label for a rib, or a vertebra as well as measurement tasks inside an anatomical structure.

7.2 Embedding Anatomy-Driven Reformations

To alleviate the tasks listed in the previous Section, we propose an integrated visualization framework enhanced by ADR views. Our framework aims to allow radiologists to detect and assess anomalies faster while reducing the risk of missing important pathologies.

The most widely practiced way to explore medical datasets is to inspect slices from standard orthogonal viewing directions (axial, coronal, sagittal) since these provide a genuinely unaltered view of the original reconstruction. For this reason, our overall concept is to use the ADR volume as an additional view for navigation and a fast overview of the structure of interest. This means we do not propose to replace original reconstructions by the ADR view, but to combine

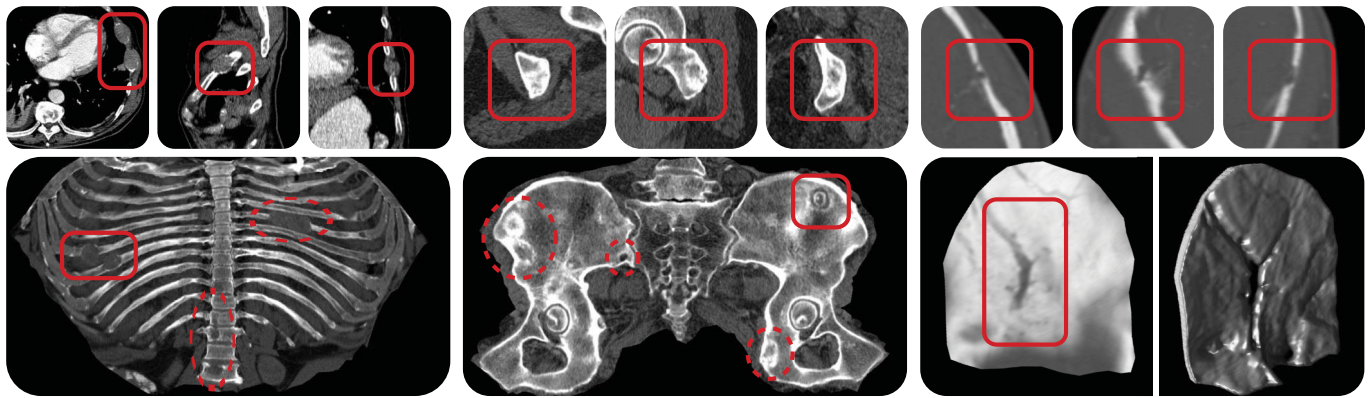


Fig. 11. Three example applications for ADR; *Left to right*: a rib cage with tumors, a pelvic bone with tumors and a potential skull fracture. The upper row shows orthogonal Reformations (axial, sagittal and coronal) and the lower row shows the center slice of the ADR volume. The lesions marked by solid boxes are shown in both reformations, whereas the dotted frames indicate additional pathologies instantly visible in the ADR view. Compared to orthogonal views, the ADR view provides a more condensed overview and it allows for an easier assessment of lesion geometry and location with respect to the anatomical structure. In case of the potential skull fracture, neither the standard views, nor the ADR is conclusive. A detail volume rendering of the inside of the reformatted skull area, however, reveals that the abnormality is in fact an emissary vein. For applications, we propose using a combination of the ADR view (detection, navigation) and original reconstructions (ultimate diagnosis).

the advantages of both: The image fidelity of the original reconstruction and the fast access and aggregated character of the ADR. Based on this approach, we propose a window layout as depicted in Figure 12 which consists of the ADR view, three linked standard views and an additional side view with a context-dependent variety of tools. For example, to perform measurements on fractures, a measurement tool can be provided to draw curves in the ADR view. After transforming the curve back to world space, very precise length measurements, that follow the anatomical geometry can be performed. As an additional tool, curves drawn in the ADR view could also be used to create CPR visualizations of rib bones or spine regions. The condensed ADR view allows doing this with fewer interactions and less slicing as compared to standard views. When inspecting ADRs, it might furthermore be advantageous to integrate information over a certain depth, similar as proposed in [34] or in [9]. All common projection algorithms such as Maximum-, Minimum-, and Mean Intensity Projections can be combined with our reformation as illustrated in Figure 10.

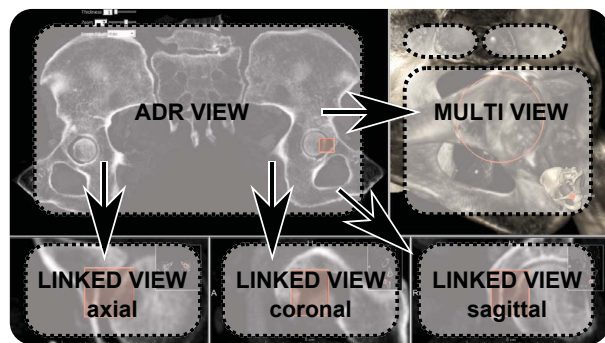


Fig. 12. Our proposal for embedding ADR in medical applications. The ADR VIEW shows the reformatted volume, supports slicing and different projections and acts as an anatomical atlas. The current viewing position and the zoom of the LINKED VIEWS is synchronized with the ADR view. The MULTI VIEW panel shows a contextual volume rendering which can be replaced by dynamic context-dependent tools.

8 CONCLUSION

We present ADR as a novel deformation-based approach to medical image reformation. The anatomy-awareness of ADR helps to facilitate many time-consuming tasks like diagnosis, navigation or annotation

and integrates seamlessly with existing visualization concepts. By explicitly minimizing deviations from local rigidity during deformation, our reformations expose very low distortions which we confirm in a diverse set of examples. To provide further control over the distribution of errors, we introduce a weighting scheme that allows increasing the faithfulness of reformations in areas of elevated importance. To demonstrate the usefulness of our method, we present several clinical applications and tasks that can increase their efficiency by introducing ADRs.

ADR is based on surface meshes that are usually placed in the center of the anatomical structure subject to investigation. This means that the quality of the final reformation depends on the precision and reliability of that surface. Since, however, our volumetric approach parameterizes a thick surrounding of the ADR surface, it is less sensitive to small errors of the input data than current cut surface-based methods. One issue is that for some shapes of ADR surfaces, the off-set surfaces may expose local or global self-intersections. In Section 4.3.1 we propose a strategy that usually removes local artifacts but for very thick and curved reformations global self-intersections cannot be completely avoided. In the variety of examples we presented in this paper, however, self-intersections were not an issue. Moreover, even if they occur, the mapping between world space and ADR space simply loses bijectivity but the approach does not break.

A limitation of the presented method is that it currently only explicitly models and minimizes the intra-layer distortion for 3 discrete layers. For current applications, however, three layers are sufficient since the corresponding anatomical structures can be parameterized by the surrounding of a surface. One way to extend the presented method would be to allow an arbitrary number of layers.

Current applications (see Section 7.1) focus on bone structures. We are, however, currently exploring different applications for our method like colon unfolding or the reformation of organs like kidneys or the heart. A restriction of our reformation approach, relevant in this respect is, that the structure the ADR surface is limited to open surfaces. This restriction is however necessary for the existence of a continuous flat embedding of a structure and does not pose problems for the proposed application to bones.

Our approach has shown promising results for oncological and trauma-related bone reading tasks and we see great potential for other applications like whole-body atlases or side-by-side comparisons. We are currently exploring semi-automatic variations of our method for ad hoc ADRs based on user sketches. Future work will focus on further developing and quantitatively evaluating our prototype application with respect to the efficiency gain for specific clinical tasks.

REFERENCES

- [1] T. Auzinger, G. Mistelbauer, I. Baclija, R. Scherthaner, A. Köchl, M. Wimmer, E. Gröller, and S. Bruckner. Vessel visualization using curved surface reformation. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2858–2867, 2013.
- [2] A. V. Bartolí, R. Wegenkittl, A. König, and E. Gröller. Nonlinear virtual colon unfolding. In *Proc. of IEEE Visualization*, pages 411–420, 2001.
- [3] M. Bhargavan, J. H. Sunshine, and B. Schepps. Too few radiologists? *American Journal of Roentgenology*, 178(5):1075–1082, 2002.
- [4] G. Chintalapani, L. M. Ellingsen, O. Sadowsky, J. L. Prince, and R. H. Taylor. Statistical atlases of bone anatomy: construction, iterative improvement and validation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 499–506. Springer, 2007.
- [5] T. F. Cootes, G. J. Edwards, C. J. Taylor, et al. Active appearance models. *IEEE Transactions on pattern analysis and machine intelligence*, 23(6):681–685, 2001.
- [6] T. F. Cootes and C. J. Taylor. Active shape modelssmart snakes. In *BMVC92*, pages 266–275. 1992.
- [7] M. De Maeseneer, J. De Mey, L. Lenchik, H. Everaert, and M. Osteaux. Helical ct of rib lesions: a pattern-based approach. *American Journal of Roentgenology*, 182(1):173–179, 2004.
- [8] T. Delamé, C. Roudet, and D. Faudot. From a medial surface to a mesh. In *Computer Graphics Forum*, volume 31, pages 1637–1646, 2012.
- [9] V. Dicken, B. Wein, H. Schubert, J.-M. Kuhnigk, S. Kra, and H.-O. Peitgen. Novel projection views for simplified reading of thorax CT scans with multiple pulmonary nodules. *Proc. of CARS*, 1256(0):59–64, 2003.
- [10] A. du Bois d’Aische, M. De Craene, S. Haker, N. Weisenfeld, C. Tempany, B. Macq, and S. K. Warfield. Improved non-rigid registration of prostate MRI. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 845–852. Springer, 2004.
- [11] N. Faraj, J.-M. Thiery, I. Bloch, N. Varsier, J. Wiart, and T. Boubekeur. Robust and scalable interactive freeform modeling of high definition medical images. In *Mesh Processing in Medical Image Analysis 2012*, pages 1–11. Springer, 2012.
- [12] M. S. Floater. Parametrization and smooth approximation of surface triangulations. *Computer aided geometric design*, 14(3):231–250, 1997.
- [13] G. Girish, K. Finlay, D. Fessell, D. Pai, Q. Dong, and D. Jamadar. Imaging review of skeletal tumors of the pelvis malignant tumors and tumor mimics. *The Scientific World Journal*, 2012, 2012.
- [14] W. Hong, X. Gu, F. Qiu, M. Jin, and A. Kaufman. Conformal virtual colon flattening. In *Proc. of the 2006 ACM symposium on Solid and physical modeling*, pages 85–93. ACM, 2006.
- [15] D. K. Jeong, K. H. Lee, B. H. Kim, K. J. Kim, Y. H. Kim, V. Bajpai, and Y. G. Shin. On-the-fly generation of multiplanar reformation images independent of CT scanner type. *Journal of digital imaging*, 21(3):306–311, 2008.
- [16] W. Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A*, 32(5):922–923, 1976.
- [17] A. Kanitsar, D. Fleischmann, R. Wegenkittl, P. Felkel, and E. Gröller. CPR - curved planar reformation. In *Proc. of IEEE Visualization*, pages 37–44, 2002.
- [18] A. Kanitsar, R. Wegenkittl, D. Fleischmann, and E. Gröller. Advanced Curved Planar Reformation: Flattening of Vascular Structures. In *Proc. of IEEE Visualization*, pages 43–50, 2003.
- [19] G. A. Lavi. Mapping the coronary arteries on a sphere in CT angiography. In *Medical Imaging 2004*, pages 285–293. International Society for Optics and Photonics, 2004.
- [20] T.-C. Lee, R. L. Kashyap, and C.-N. Chu. Building skeleton models via 3-d medial surface axis thinning algorithms. *CVGIP: Graphical Models and Image Processing*, 56(6):462–478, 1994.
- [21] Z. Levi and D. Levin. Shape deformation via interior RBF. *IEEE Transactions on Visualization and Computer Graphics*, 20(7):1062–1075, July 2014.
- [22] B. Lévy, S. Petitjean, N. Ray, and J. Mailliot. Least squares conformal maps for automatic texture atlas generation. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 362–371. ACM, 2002.
- [23] L. Liu, L. Zhang, Y. Xu, C. Gotsman, and S. J. Gortler. A local/global approach to mesh parameterization. In *Computer Graphics Forum*, volume 27, pages 1495–1504, 2008.
- [24] S. Makram-Ebeid and O. Somphone. Non-rigid image registration using a hierarchical partition of unity finite element method. In *Proc. of IEEE Computer Vision*, pages 1–8. IEEE, 2007.
- [25] L. S. Medina. Three-dimensional CT maximum intensity projections of the calvaria: a new approach for diagnosis of craniosynostosis and fractures. *American journal of neuroradiology*, 21(10):1951–1954, 2000.
- [26] G. Mistelbauer, A. Varchola, H. Bouzari, J. Starinsky, A. Köchl, R. Scherthaner, D. Fleischmann, M. E. Gröller, and M. Srámek. Centerline reformations of complex vascular structures. In *IEEE PacificVis*, pages 233–240, 2012.
- [27] A. Nealen, T. Igarashi, O. Sorkine, and M. Alexa. Laplacian mesh optimization. In *Proc. of GRAPHITE*, pages 381–389. ACM, 2006.
- [28] D. S. Paik, C. F. Beaulieu, R. B. Jeffrey Jr, C. A. Karadi, and S. Napel. Visualization modes for CT colonography using cylindrical and planar map projections. *Journal of computer assisted tomography*, 24(2):179–188, 2000.
- [29] D. Pavić and L. Kobbelt. High-resolution volumetric computation of off-set surfaces with feature preservation. In *Computer Graphics Forum*, volume 27, pages 165–174, 2008.
- [30] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics*, 2(1):15–36, 1993.
- [31] B. Reiner. Automating radiologist workflow, part 2: hands-free navigation. *Journal of the American College of Radiology*, 5(11):1137–1141, 2008.
- [32] B. I. Reiner, E. L. Siegel, F. J. Hooper, S. Pomerantz, A. Dahlke, and D. Rallis. Radiologists’ productivity in the interpretation of CT scans: a comparison of pacs with conventional film. *American Journal of Roentgenology*, 176(4):861–864, 2001.
- [33] C. Rieder, A. Weihusen, C. Schumann, S. Zidowitz, and H.-O. Peitgen. Visual support for interactive post-interventional assessment of radiofrequency ablation therapy. 29(3):1093–1102, 2010.
- [34] H. Ringl, R. E. Scherthaner, G. Schueller, C. Balassy, D. Kienzl, A. Botosaneanu, M. Weber, C. Czerny, S. Hajdu, T. Mang, et al. The skull unfolded: a cranial CT visualization algorithm for fast and easy detection of skull fractures. *Radiology*, 255(2):553, 2010.
- [35] J. R. Rossignac and A. A. Requicha. Offsetting operations in solid modelling. *Computer Aided Geometric Design*, 3(2):129–148, 1986.
- [36] L. Saroul, O. Figueiredo, and R.-D. Hersch. Distance preserving flattening of surface sections. *IEEE Transactions on Visualization and Computer Graphics*, 12(1):26–35, Jan 2006.
- [37] L. Saroul, S. Gerlach, and R. D. Hersch. Exploring curved anatomic structures with surface sections. In *Proc. of the IEEE Visualization*, pages 5–12, 2003.
- [38] H. Seim, D. Kainmueller, M. Heller, H. Lamecker, S. Zachow, and H.-C. Hege. Automatic segmentation of the pelvic bones from CT data based on a statistical shape model. In *EG Workshop on Visual Computing in Biology and Medicine*, pages 93–100, 2008.
- [39] A. Shamir. A survey on mesh segmentation techniques. *Computer Graphics Forum*, 27(6):1539–1556, 2008.
- [40] O. Sorkine and M. Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, volume 4, pages 109–116, 2007.
- [41] M. Termeer, J. Bescos, M. Breeuwer, A. Vilanova, F. Gerritsen, and E. Gröller. Covicad: Comprehensive visualization of coronary artery disease. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1632–1639, Nov 2007.
- [42] D. Williams, S. Grimm, E. Coto, A. Roudsari, and H. Hatzakis. Volumetric curved planar reformation for virtual endoscopy. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):109–119, Jan 2008.
- [43] D. Wu, D. Liu, Z. Puskas, C. Lu, A. Wimmer, C. Tietjen, G. Soza, and S. K. Zhou. A learning based deformable template matching method for automatic rib centerline extraction and labeling in CT images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 980–987, 2012.
- [44] W. Zeng, Y. Zeng, Y. Wang, X. Yin, X. Gu, and D. Samaras. 3d non-rigid surface matching and registration based on holomorphic differentials. In *Computer Vision-ECCV 2008*, pages 1–14. Springer, 2008.
- [45] S. Zhang, A. Nealen, and D. Metaxas. Skeleton based as-rigid-as-possible volume modeling. In *Eurographics 2010-Short Papers*, pages 21–24, 2010.
- [46] M. Zollhöfer, E. Sert, G. Greiner, and J. Süßmuth. GPU based arap deformation using volumetric lattices. In *Eurographics 2012-Short Papers*, pages 85–88. The Eurographics Association, 2012.
- [47] G. Zou, J. Hua, and O. Muzik. Non-rigid surface registration using spherical thin-plate splines. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 367–374. 2007.