



# Skuller: A volumetric shape registration algorithm for modeling skull deformities



Yusuf Sahillioğlu<sup>a,\*</sup>, Ladislav Kavan<sup>b</sup>

<sup>a</sup> Computer Engineering Department, Middle East Technical University, Turkey

<sup>b</sup> Computer & Information Science Department, University of Pennsylvania, USA

## ARTICLE INFO

### Article history:

Received 17 March 2014

Received in revised form 22 December 2014

Accepted 11 March 2015

Available online 8 April 2015

### Keywords:

Volumetric shape registration

Deformation

High-resolution model

## ABSTRACT

We present an algorithm for volumetric registration of 3D solid shapes. In comparison to previous work on image based registration, our technique achieves higher efficiency by leveraging a template tetrahedral mesh. In contrast to point- and surface-based registration techniques, our method better captures volumetric nature of the data, such as bone thickness. We apply our algorithm to study pathological skull deformities caused by a particular condition, i.e., craniosynostosis. The input to our system is a pair of volumetric 3D shapes: a tetrahedral mesh and a voxelized object represented by a set of voxel cells segmented from computed tomography (CT) scans. Our general framework first performs a global registration and then launches a novel elastic registration process that uses as much volumetric information as possible while deforming the generic template tetrahedral mesh of a healthy human skull towards the underlying geometry of the voxel cells. Both data are high-resolution and differ by large non-rigid deformations. Our fully-automatic solution is fast and accurate, as compared with the state of the arts from the reconstruction and medical image registration fields. We use the resulting registration to match the ground-truth surfaces extracted from the medical data as well as to quantify the severity of the anatomical deformity.

© 2015 Elsevier B.V. All rights reserved.

## 1. Introduction

Shape registration is an important problem in computer graphics with numerous applications from shape completion to animation reconstruction. Registration algorithms aim to deform a source shape into a target shape while optimizing a prescribed criteria, which generally involves spatial closeness between the final pose and the target. Registration algorithms for shapes represented as point clouds and surfaces have been well-studied (Tam et al., 2013), but this problem is less well-understood for registering volumetric source and target shapes. Our approach is, to our knowledge, the first one to address this version of the problem, and differs ideologically from volumetric registration in medical imaging where displacement vectors are optimized over the whole image volume, not over the 3D solid shapes.

Explicit volumetric representations of solid 3D objects that are common in computer graphics are preferable to thin shell surfaces for the purposes of realistic animation or physically-based elastic simulation. One may also use volumetric models to describe anatomical structures in 3D medical imaging. With our method

we explicitly represent one such structure, namely the human skull, as a volumetric tetrahedral mesh which facilitates operations such as shape analysis and visualization. With surface representations, on the other hand, it is difficult to account for the volumetric thickness of the skull bones.

The main motivation for our volumetric registration algorithm is that we introduce a general framework to this new volumetric shape registration. Namely, we adapt volumetric shape registration techniques from computer graphics domain to the medical imaging field, and introduce several novel concepts to make this process work well. Our medical objective is craniosynostosis, a condition caused by premature fusion of individual bone plates in the skull, which results in a deformed head shape in children. Unfortunately, quantitative objective measure of the severity of this condition is lacking. We believe that if we provide a *registered* deformable model of the patient to a doctor equipped with pre-defined tools such as a deformation rig, then he can better assess the condition (Section 7.2.2). Such a *registered* model can also be used as a transfer intermediate to establish dense correspondences (Section 7.2.1).

In this paper we present a novel fully-automatic non-rigid volumetric shape registration framework that is fast enough to handle high-resolution input data in minutes and is accurate enough to

\* Corresponding author.

capture highly distorted skull models of actual craniosynostotic patients. The input to our system is explicitly volumetric shapes, namely, a high-quality generic tetrahedral template mesh to be deformed towards the segmented skull voxels of computed tomography (CT) scans of a patient. To achieve the registration, we first establish an initial alignment between the data sets by a global rigid transformation and uniform scale. Non-rigid transformation is then computed by a quadratic deformation model that leverages as much volumetric information as possible while matching the mesh tetrahedra and voxel cells. Note that, although some FEM-based elastic registrations take advantage of volumetric tetrahedral meshes for a more realistic deformation, these meshes conform to non-volumetric point clouds. Our tetrahedral mesh, on the other hand, takes further advantage of the volumetricity by conforming to a volumetric shape comprised of skull voxels. Thanks to the simple closed form of our deformation model, all our algorithm boils down to iteratively solving a sparse linear system, which makes it quite fast. Our accurate registration leads to two applications, shape matching and anatomical deformity quantification.

The source code and the executables for the method that we present in this paper are publicly available in <http://www.ceng.metu.edu.tr/~ys/pubs>.

## 2. Related work

The transformations admitted by shape registration methods for point clouds and surfaces have evolved step by step from rigid to non-rigid. The prominent rigid registration algorithm is Iterative Closest Points (ICP) algorithm (Besl and McKay, 1992) that alternates between correspondence and rigid transformation optimization. Various stages of ICP have been improved such as matching point selection and error function construction (Rusinkiewicz and Levoy, 2001). ICP may, however, easily converge to a local minimum of the error function if the initial positioning of the shape pair is not sufficiently close. This problem is addressed by providing a good initialization based on matching of local descriptors that are invariant under rigid transforms (Gelfand et al., 2005; Huang et al., 2006) or based on a randomized approach that is more robust to noise and outliers (Aiger et al., 2008). As a faster alternative we use multiple initialization candidates generated by PCA-based pose normalization (Section 4), a side contribution of this research.

For registration methods that support non-rigid transformations there is an important issue of maintaining data regularity as displacement of each point may arbitrarily deviate from the neighboring ones if precautions are not taken. These arbitrary transformations eventually lead to distortions, e.g., degenerate triangles on surface. Smoothness is attained in Chui and Rangarajan (2003) by specializing the non-rigid transformation function to thin-plate splines, whereas Hahnel et al. (2003) uses multiple landmark measurements as boundary conditions in the optimization. A combination of the classical Laplacian smoothing and surface fairing technique (Taubin, 1995) is employed in Sahillioglu and Yemez (2010) as a smoothing operator. Template-based method of Allen et al. (2003) explicitly minimizes the difference of transformations between every two points that are adjacent in the deforming mesh along with other objectives controlling the spatial closeness of automatically and manually matched points. Li et al. (2013) proposes yet another simultaneous optimization regulated by a deformation graph model (Sumner et al., 2007).

Isometric deformations are a subset of non-rigid transformations after which pairwise geodesic distances between the shapes are preserved. If one may assume isometry, then the correspondence computation that guides the registration process

can be improved significantly by injecting second-order terms in the optimization in addition to or instead of the first-order point-wise terms to be compared (Sahillioglu and Yemez, 2012; Sahillioglu and Yemez, 2013). In particular, Zhang et al. (2008) prunes the combinatorial search space of anchor point correspondences according to some criteria based on local shape similarity and geodesic consistency and performs the deformation with these matches as constraints. Huang et al. (2008) employs a similar geodesic-based pruning for the correspondences that guide the deformation in which rigid transformations are applied to the group of points that can move rigidly. Similarly, Chang and Zwicker (2009) clusters groups of points moving together to solve for skinning weights. Geodesic consistency is enforced in Tevs et al. (2009) with randomized feature matching and in template-based Anguelov et al. (2004) with an optimization of a joint probabilistic model over all point-to-point correspondences between two shapes. Both methods can handle large isometric deformations given sufficient computation time.

For registration of anatomical structures in 3D medical images (McInerney and Terzopoulos, 1996; Holden, 2008; Markelj et al., 2012; Sotiras et al., 2013), there are methods based on elastic surface models and volumetric deformation models. The former geometrically deforms an initial spherical surface (Miller et al., 1991; McInerney and Terzopoulos, 1995) or the CT point cloud (Feldmar et al., 1996) or a finite element model (FEM) (Gee et al., 1994; Ferrant et al., 2001; Schnabel et al., 2003; Cash et al., 2005; Bucki et al., 2010) until it conforms to the target surface in the CT data. The latter, on the other hand, formulates the deformation as a volumetric transformation based on free-form deformations (Szeliski and Lavallee, 1996; Christensen et al., 1997; Rueckert et al., 1999; Rohlfing et al., 2003; Ou et al., 2011; Pszczolkowski et al., 2012) or volumetric radial basis functions (Pighin et al., 1998; Kahler et al., 2002; Timoner, 2003) or diffeomorphism (Beg et al., 2005). Volumetric models, unless designed in coarse-to-fine fashion, have significantly higher computation costs than elastic models as they represent deformations over a dense 3D region instead of over a manifold. This representation may also degrade accuracy unless the user marks the critical parts whose surfaces are close spatially but far geodesically, such as eye-holes. Surface and point-based elastic methods, on the other hand, obviously fail to capture volumetric thicknesses.

FEM-based volumetric registration techniques can be employed to efficiently create a subject-specific FE model by morphing a pre-defined generic mesh onto the target organ. The generic mesh, which resembles an average model from the target organ population, is typically constructed manually or semi-manually to ensure mesh quality. This mesh can represent complex geometries, define anisotropic material properties for constituent elements, and also address specific boundary conditions, which in turn make the FEM-based approaches quite popular. The shortcoming of the FEM-based registration, on the other hand, is the degradation in mesh quality, e.g., irregular elements and the introduction of inverted elements with negative volumes, which can occur because of excessive mesh distortion during the morphing. While some methods alleviate this issue with algorithmically complicated and time consuming mesh repair operations applied after the mesh is transformed, our method, which is a simple variant of the FEM method, minimizes mesh distortion during the registration by employing volumetric regularization terms.

Although we see tetrahedral source meshes in some FEM-based elastic registrations, they are controlled by non-volumetric point clouds on the target side of the pipeline, unlike our algorithm that employs an explicitly volumetric space of target voxel cells. Similarly, elastic free-form deformation methods maintain an embedded 2D manifold inside a deforming volumetric cage, which is again attracted to a non-volumetric point cloud. In the view of

the above discussion, we claim that our algorithm is the first one to deform an explicit volumetric *shape*, such as a tetrahedral mesh, towards another explicit volumetric *shape*, such as a grid of voxel cells. Addressing the registration problem at shape-level brings extreme efficiency without losing much accuracy compared to the conventional volumetric models that work on the whole or segmented medical images, as shown in Section 7.3. Finally, atlas-based registration methods in the medical field aim to evolve template anatomic atlases to images (Chen and Medioni, 1991; Yao and Taylor, 2000; D’Agostino et al., 2004; Chintalapani et al., 2007; Ellingsen et al., 2010). In general the templates used are other images, but in this work a tetrahedral mesh of the organ is used, which is potentially more promising to serve for both visualization and fast modeling.

Instead of registering the deformable source model to the reference target model, one may directly reconstruct the latter without using deformations at all. Multiple close layers of points in CT scans that we address in this work fail the Voronoi-based reconstruction approaches (Amenta et al., 2001; Dey et al., 2010). This setting also makes consistent normal computation difficult and inevitably unstable, which in turn degrades performances of signed distance field based methods (Hoppe et al., 1992; Frisken and Rockwood, 2000; Kobbelt et al., 2001) and Poisson reconstruction (Kazhdan and Hoppe, 2013). Yet even with the assumed normals, Poisson fits only one surface to the volumetric data whose underlying geometry may well include multiple close-by outer and inner surfaces, as in the case of skull bones. A problem common to reconstruction methods is that they become severely under-constrained when the amount of missing data is large, leading to poor results. Tagliassachi et al. (2011) addresses this issue with a computationally expensive volume-aware approach using view vectors at each point, which is again hard to achieve in CT scans. Alleviating these problems is possible by using the continuous intensity information in medical data with the standard Marching Cubes surface extraction method (Lorensen and Cline, 1987). This intensity-based ground-truth reconstruction can be made even more valuable by matching the extracted surfaces with each other, as we achieved in Section 7.2.1.

### 3. Problem statement and overview

Our goal is to register volumetric data; in particular, a high-quality generic template mesh and a patient-specific CT scan. Our pipeline is depicted in Fig. 1.

The input data sets are first pre-processed. We first segment the desired 3D skull voxels from the CT scans using a simple thresholding based on intensity values. Our generic skull template, on the other hand, is obtained from Zygote (2013) as a triangular mesh and converted to a tetrahedral mesh (Jacobson et al., 2013). We denote these volumetric geometries as  $\mathcal{T}$  for tetrahedral template mesh, and  $\mathcal{S}$  for CT scan skull voxels, as visualized in Fig. 1.

### 4. Initial alignment

Our template mesh is given in an arbitrary scale, position, and orientation. Our first task is therefore to provide a global alignment by allowing only translation, rotation, and uniform scale. We apply a standard ICP algorithm (Besl and Mckay, 1992). The initial guess for ICP is calculated using PCA alignment (Kazhdan, 2007). After mean alignment (global translation), we compute the principal components of the covariance matrices for both shapes. For this purpose, the tetrahedral mesh  $\mathcal{T}$  is considered just as 3D vertices, and similarly CT skull voxels  $\mathcal{S}$  is considered just as 3D center points of each segmented voxel. We then consider all of the 4 possible rotations (mapping one basis of principal components to the other one). We execute ICP on each of these four possibilities, and choose the one with the lowest error (Fig. 2 – left). We generalize the classical rigid ICP by allowing for uniform scale, because the scale of our template mesh is arbitrary. Uniform scale is commutative with rotation and therefore only a minimal change of the ICP algorithm is required. Namely, we jointly optimize the uniform scale factor  $u$  and translation  $\mathbf{d}$  to be applied to template points rotated by ICP. Let  $\mathbf{p}$  denote these rotated points in correspondence with points  $\mathbf{q}$  from the target shape. We then minimize the error of the current transform  $f(\mathbf{p}) = \sum_i \|u\mathbf{p}_i + \mathbf{d} - \mathbf{q}_i\|^2$ . This novel modification, as a side contribution of this research, results in a better initial alignment, as shown in Fig. 2 – right.

### 5. Volumetric registration

Let us denote the rest-pose vertex positions of our template mesh (after global registration) as  $\mathbf{t}_0 \in \mathbb{R}^{3n \times 1}$  where  $n$  is the number of mesh vertices, and 3 elements are stored for each vertex, namely the  $x$ -,  $y$ -, and  $z$ -coordinates. Similarly let  $\mathbf{s} \in \mathbb{R}^{3m \times 1}$  denote the coordinates of the centers of the  $m$  input voxels.

We want to find deformed vertices  $\mathbf{t}$  of the template mesh such that they match  $\mathbf{s}$  as closely as possible without deviating from the rest pose  $\mathbf{t}_0$  more than is necessary. Therefore, we formulate an energy term that balances the accuracy of the match between  $\mathbf{t}$  and  $\mathbf{s}$  with deformation energy that penalizes deviation of  $\mathbf{t}$  from the rest pose  $\mathbf{t}_0$ . To avoid overly large steps that could result in unreliable correspondences, we also introduce a Tikhonov regularization term that penalizes too large motion from the previous step.

Our deformation model seeks  $\mathbf{t} \in \mathbb{R}^{3n \times 1}$  that minimizes the following energy:

$$E_{\text{def}}(\mathbf{t}) = E_{\text{corr}}(\mathbf{t}) + \alpha E_{\text{Dirichlet}}(\mathbf{t}) + \beta E_{\text{Tikhonov}}(\mathbf{t}) \quad (1)$$

The first energy term  $E_{\text{corr}}$  concerns correspondence between the transformed mesh vertices  $\mathbf{t}$  and the fixed voxels  $\mathbf{s}$ . We define our correspondences using a binary matrix  $\mathbf{P}$  of size  $n \times m$  (i.e., each

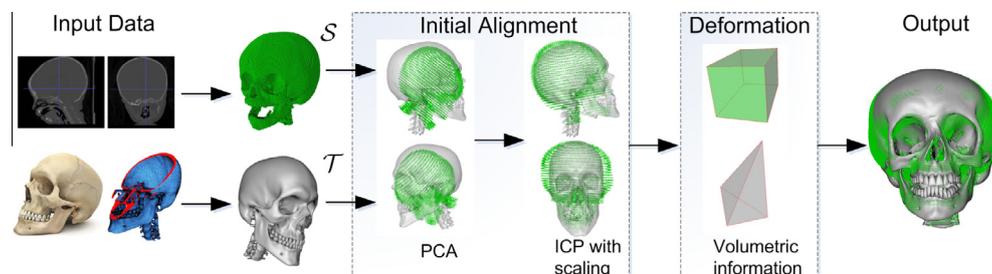
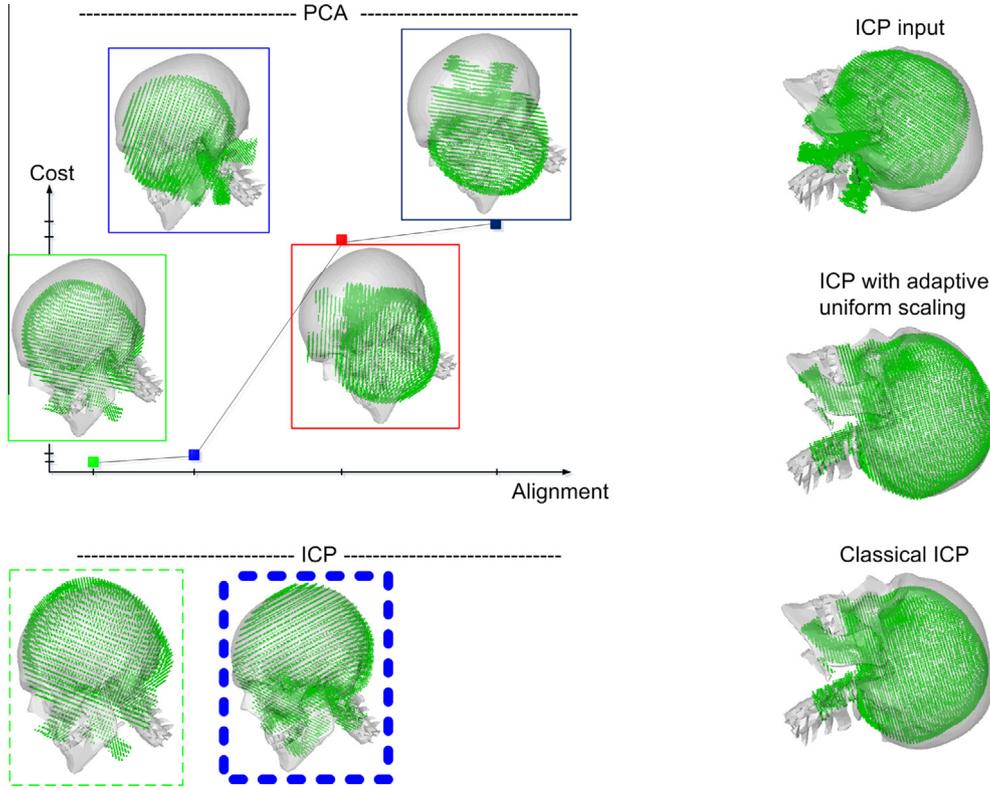


Fig. 1. The overview of our volumetric registration framework. We show multiple views for both steps of the initial alignment.



**Fig. 2.** (Left) Out of four PCA-based alignment candidates that initialize ICP, the correct blue one results in the closest ICP alignment although it has a higher PCA alignment cost. (Right) ICP result with (middle) and without (bottom) our adaptive uniform scaling modification. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

element is either 0 or 1), such that each row contains only one “1” and the rest are zeros. The fact that  $\tilde{P}_{ij} = 1$  means that voxel number  $j$  is in correspondence with mesh vertex  $i$ , i.e.,  $i$  is the vertex closest to voxel  $j$ . We define  $\mathbf{P} := \tilde{\mathbf{P}} \otimes \mathbf{I}_3$ , where  $\mathbf{I}_3$  denotes  $3 \times 3$  identity and  $\otimes$  is the Kronecker product. The purpose of Kronecker extension is to be able to select 3 elements, namely the  $x$ -,  $y$ -, and  $z$ -coordinates, of the corresponding voxel from  $\mathbf{s}$ .

For a given matrix  $\mathbf{P}$ , the correspondence energy is then defined as follows.

$$E_{\text{corr}}(\mathbf{t}) = \|\mathbf{t} - \mathbf{P}\mathbf{s}\|^2 \quad (2)$$

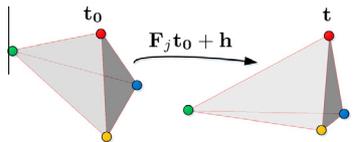
As a deformation energy, we found it is sufficient to use the (volumetric) Dirichlet energy of the displacements:

$$E_{\text{Dirichlet}}(\mathbf{t}) = (\mathbf{t} - \mathbf{t}_0)^T \mathbf{L}(\mathbf{t} - \mathbf{t}_0) \quad (3)$$

where  $\mathbf{L}$  is the mesh Laplacian. We calculate  $\mathbf{L}$  as

$$\mathbf{L} = \sum_{j \in \mathcal{T}} \lambda_j \mathbf{G}_j^T \mathbf{G}_j \quad (4)$$

where  $\lambda_j$  is the volume of the tetrahedron  $j$  at rest-pose and  $\mathbf{G}_j$  is a constant matrix to obtain the deformation gradient  $\mathbf{F}_j$  of tetrahedron  $j$  (Fig. 3). Intuitively,  $\mathbf{G}_j \in \mathbb{R}^{9 \times 3n}$  extracts the vectorized version of the desired gradient  $\mathbf{f}_j = \text{vec}(\mathbf{F}_j) \in \mathbb{R}^{9 \times 1}$  when multiplied with the whole set of vertices  $(\mathbf{t} - \mathbf{t}_0) \in \mathbb{R}^{3n \times 1}$ , i.e.,  $\mathbf{f}_j = \mathbf{G}_j(\mathbf{t} - \mathbf{t}_0)$ .



**Fig. 3.** Deformation gradient  $\mathbf{F}_j$  mapping the rest-pose edges of tetrahedron  $j$  to their deformed configuration minus the translation  $\mathbf{h}$ .

Additionally, we employ a Tikhonov regularization term that controls the amount of displacement from the *previous* step (as opposed to the deformation energy which penalizes deformation from the rest pose). This term is defined as:

$$E_{\text{Tikhonov}}(\mathbf{t}) = \|\mathbf{t} - \mathbf{t}_{\text{prev}}\|^2 \quad (5)$$

where  $\mathbf{t}_{\text{prev}}$  is the state of the previous iteration.

Differentiating our energy with respect to  $\mathbf{t}$  gives:

$$\frac{\partial E_{\text{def}}}{\partial \mathbf{t}} = \frac{\partial E_{\text{corr}}}{\partial \mathbf{t}} + \alpha \frac{\partial E_{\text{Dirichlet}}}{\partial \mathbf{t}} + \beta \frac{\partial E_{\text{Tikhonov}}}{\partial \mathbf{t}} \quad (6)$$

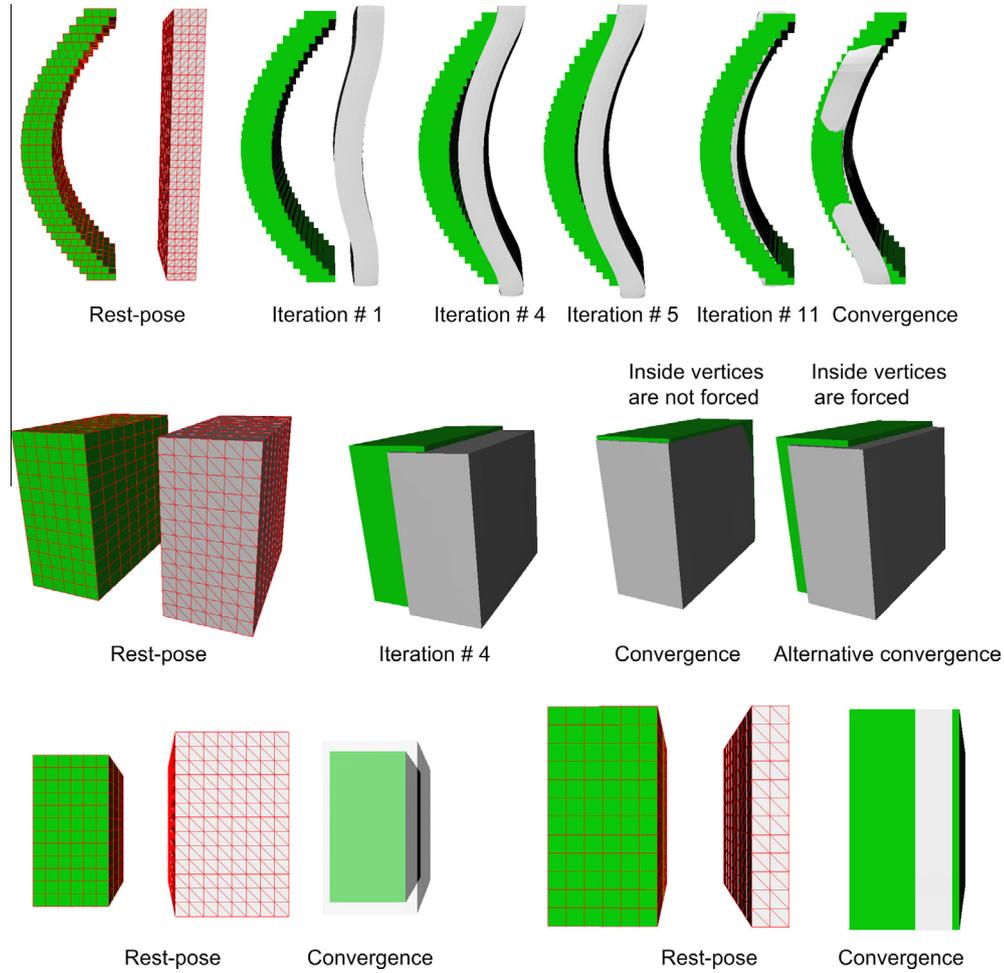
$$= 2(\mathbf{t} - \mathbf{P}\mathbf{s} + \alpha \mathbf{L}(\mathbf{t} - \mathbf{t}_0) + \beta(\mathbf{t} - \mathbf{t}_{\text{prev}})) \quad (7)$$

Setting this derivative to zero, we obtain a linear system:

$$(\alpha \mathbf{L} + (\beta + 1)\mathbf{I})\mathbf{t} = \mathbf{P}\mathbf{s} + \alpha \mathbf{L}\mathbf{t}_0 + \beta \mathbf{t}_{\text{prev}} \quad (8)$$

which we solve using a sparse direct solver.

There are several key concepts necessary to make this process work well for deformable volumetric shape registration. First, we start with an initial value of  $\alpha$  and  $\beta$ , set to conservatively high values, namely 1 for each. We then iterate the following two steps: (1) establish the correspondence matrix  $\mathbf{P}$  by calculating closest points, and (2) minimize our energy by solving the linear system above. When the linear system solve results in little or no progress, i.e., the resulting  $\mathbf{t}$  is very close to  $\mathbf{t}_{\text{prev}}$ , we reduce the strength of the regularization terms by dividing both  $\alpha$  and  $\beta$  by 2. The intuition behind using initially high  $\alpha$  and  $\beta$  values is that the closest-point correspondences in  $\mathbf{P}$  is not very reliable when the transforming  $\mathbf{t}$  is far away from the fixed target  $\mathbf{s}$  at the first iterations. In this case, we should keep the regularization terms powerful enough to prevent a potential damage from the correspondence term. As  $\mathbf{t}$  comes closer to  $\mathbf{s}$  in subsequent iterations,  $\mathbf{P}$  gets more reliable and consequently we start to rely more on the



**Fig. 4.** (Top row)  $E_{\text{Dirichlet}}$  prevents volume shrinkage as  $E_{\text{corr}}$  is attracting gray  $\mathbf{t}$  towards green  $\mathbf{s}$  (Iteration # 1). Little displacement between intermediate Iterations # 4 and 5 detected by  $E_{\text{Tikhonov}}$  triggers adaptive weighting which enables further penetration towards a concavity. (Middle row) Inside vertices that are not forced to match are pushed from back by the outside vertices via regularization effect which results in full penetration at convergence; if forced then they resist the push backs resulting in partial penetration at an alternative convergence. (Bottom row) A thinner  $\mathbf{t}$  does not wander inside the thick volume  $\mathbf{s}$  and stays put once it is fully inside (right). A thicker  $\mathbf{t}$  covers the whole  $\mathbf{s}$  (left).

correspondence term in order to progressively capture finer and finer details of the target shape  $S$  (Fig. 4 – top row).

Another important technique is to disable the correspondence term  $E_{\text{corr}}$  for mesh vertices which are already inside  $S$ . We achieve this by using the following modified correspondence term

$$E_{\text{corr}}(\mathbf{t}) = (\mathbf{t} - \mathbf{Ps})^T \mathbf{W}(\mathbf{t} - \mathbf{Ps}) \quad (9)$$

where  $\mathbf{W}$  is a diagonal matrix, such that  $W_{i,i} = 0$  if the vertex  $i$  is already inside  $S$  and  $W_{i,i} = 1$  otherwise. This technique is based on the observation that a vertex that is already inside should not be affected by the correspondence term any more as it is not possible to establish meaningful correspondences in the interior of the mesh – instead, it should be free to move as dictated by the regularization terms and vertices that are outside of  $S$ , so that the outside vertices find their way in as well. The benefit of this idea is demonstrated in Fig. 4 – middle row along with two different thickness settings of the input (Fig. 4 – bottom row).

Note that a fourth quadratic term that would measure the squared distances between user-specified landmark points could also be incorporated into our energy if user-assisted registration was desired. However, we opt for a fully-automatic framework, which works well with our data.

## 6. Computational complexity

Our initial alignment computation is dominated by closest point searches that both exist in the PCA and ICP stages with a time complexity of  $O(n \log m)$  using kd-trees. An iteration of our deformation model first takes  $O(n)$  time for the detection of inside vertices using voxel neighborhoods, and then performs closest points search for  $E_{\text{corr}}$  in  $O(n \log m)$  time, and then finally solves a sparse linear system with the  $O(n)$ -time Cholesky factorization of a coefficient matrix whose nonzero values appear only for the vertex pairs connected by an edge in the sparse tetrahedral mesh. The number of deformation iterations is independent of  $n$  and  $m$ , and happens to be about 20 in all our runs.

## 7. Results

We have tested our algorithm on five CT scans that belong to actual 6–12 months old craniosynostotic patients who complain from sagittal synostosis, metopic synostosis, and coronal synostoses. The variety of the underlying skull shapes is sufficient to conduct the preliminary experiments of our volumetric shape registration framework. Images are taken using General Electrics LightSpeed Ultra and Philips Brilliance 64 scanners, where the axial

in-plane pixel size is ranged between 0.28 and 0.45 mm, and the axial spacing is ranged between 0.31 and 4.5 mm. The resolutions of both the CT scans and the template mesh to be deformed are challenging enough to verify the tractability of our approach. CT scans represented by  $\{S_j, j = 0, \dots, 4\}$  have approximately 190 K skull voxels (out of the total  $512 \times 512 \times 370$ ) whereas our template model consists of 160 K vertices connected by 560 K tetrahedra (Fig. 1 – second column).

In addition to the visual evidence of our performance provided through Figs. 5–7, we also demonstrate our advantage over a state of the art surface reconstruction method (Kazhdan and Hoppe, 2013) with regards to capturing bone thickness (Fig. 8). Based on our registration result, we then design two applications, each with many potential benefits such as attribute transfer (Section 7.2.1) and decision making before or during surgeries (Section 7.2.2). We finally perform quantitative (Tables 1 and 2) and visual (Fig. 11) comparisons with state of the art medical image registration algorithms. Please see the accompanying video for more results.

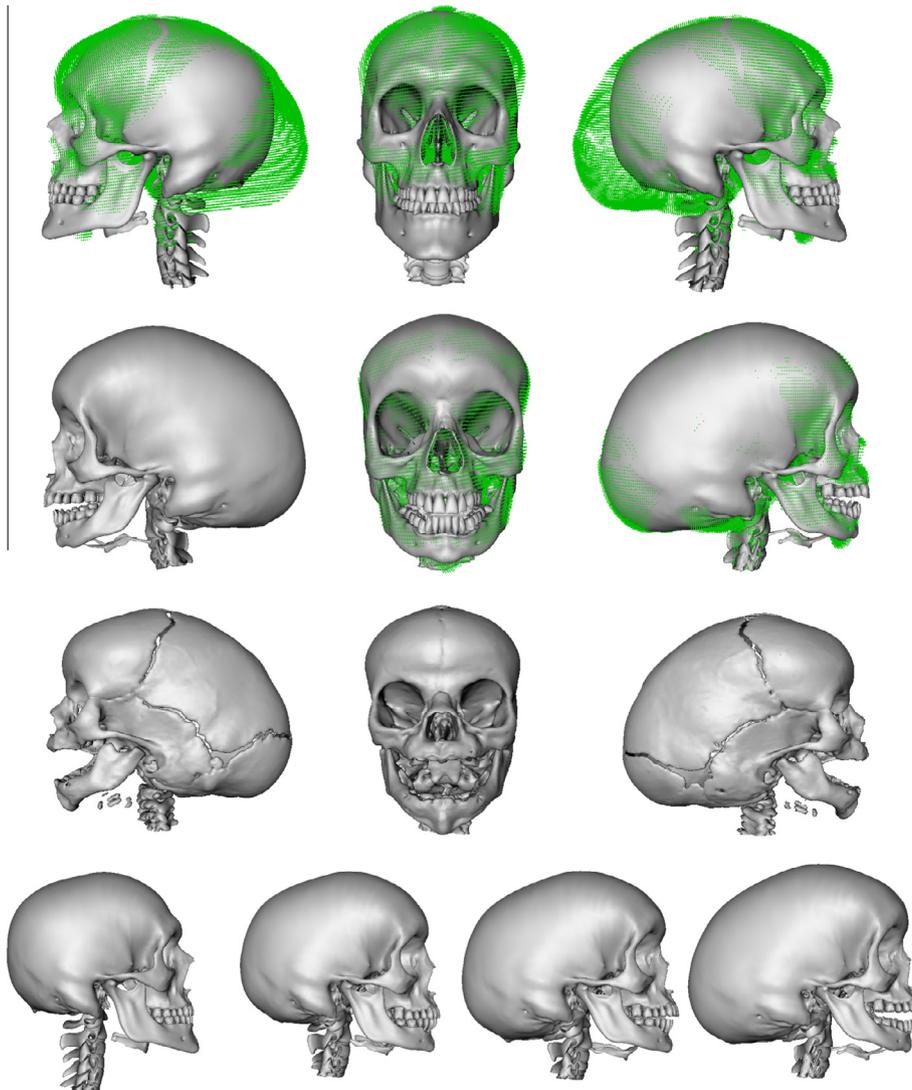
In order to enable an easier assessment of the results, we add the corresponding ground-truth surfaces extracted from the medical data (Lorensen and Cline, 1987) to all the relevant figures,

namely Figs. 5, 7 and 11. These surfaces and the ones in Fig. 9 demonstrate the finest details, such as the cracks between the bone plates, which are lacking in our resulting models. This issue, as well as the removal of the extra template parts that do not exist in the scans, such as teeth, can be handled as a future work by incorporating a remeshing operation to our system. We, however, show that the registration results, despite the missing finest details and the extra parts, can still be very useful in shape matching (Section 7.2.1) and anatomical deformity quantification (Section 7.2.2) applications.

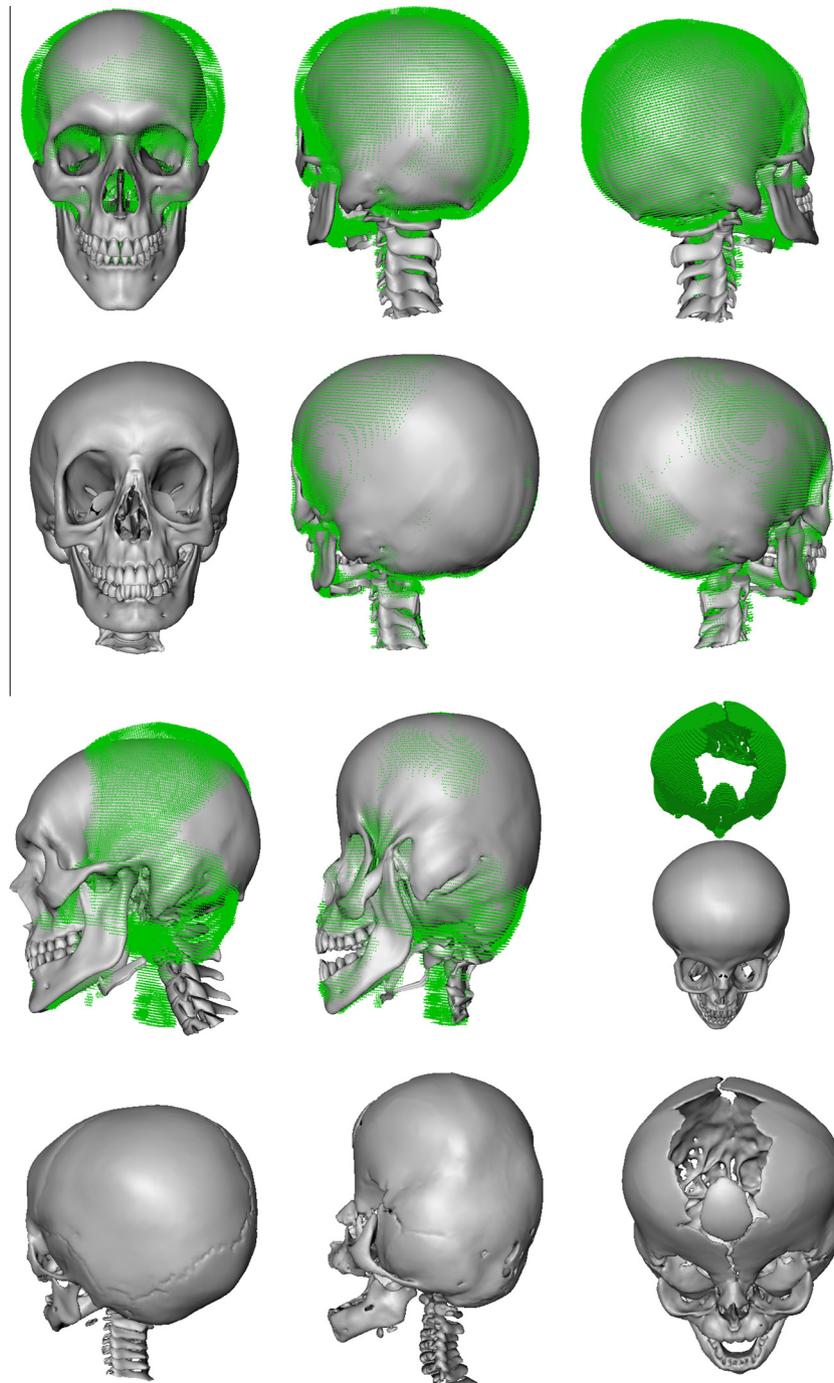
### 7.1. Visual performance and timing

In Fig. 5, we see a severely distorted skull  $S_1$  being registered with our high-quality template mesh. In addition to the global capture of the shape, note the detail capture such as the growth of the eyeholes on the template. Note also that, the back of the skulls after global alignment (Fig. 5 – first row) are still spatially apart before entering the non-rigid registration, and therefore resemble the demonstrative cases in Fig. 4.

Resulting registrations of the other scans  $S_0$  and  $S_2$  with the template is shown in Fig. 6, where we also show the robustness



**Fig. 5.** (First row) Initial alignment between our gray template mesh and green CT scan  $S_1$ . (Second row) Resulting registered model based on this initialization. (Third row) Underlying ground-truth surfaces of the CT scans. (Fourth row) Deformation process is visualized. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



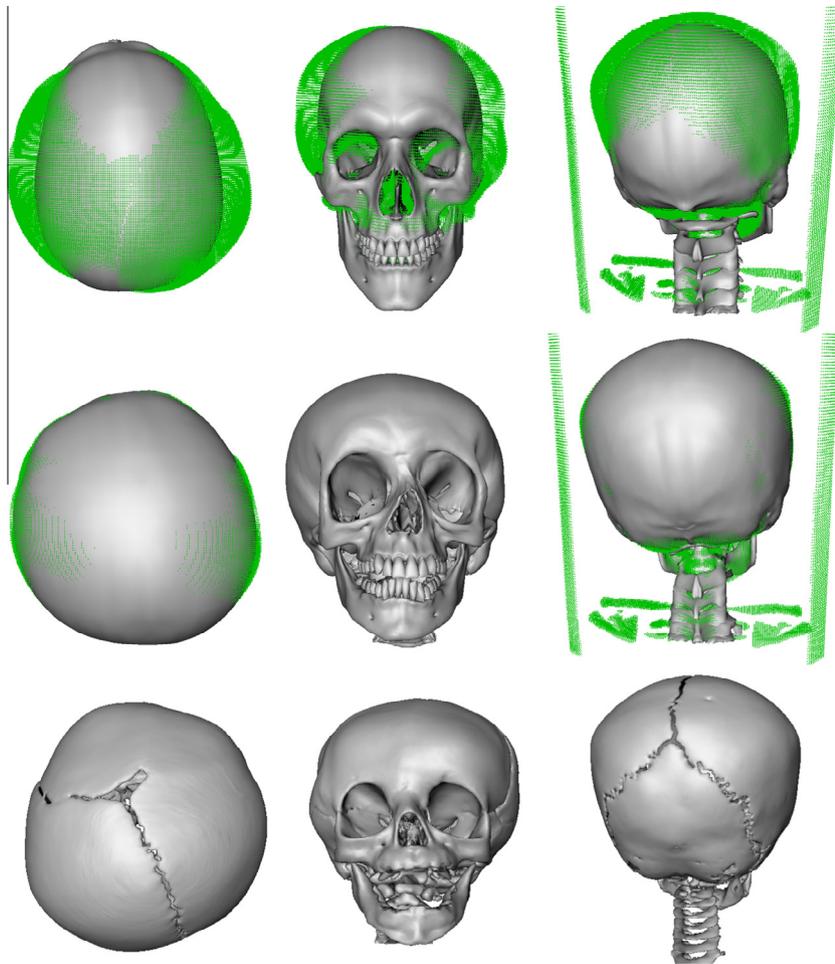
**Fig. 6.** (First row) Initial alignment between the template mesh  $\mathcal{T}$  and  $S_0$ . (Second row) Resulting registered deformable model based on this initialization. (Third row) Initial (left) and final (middle) alignment for  $S_2$ , where the hole on top of  $S_2$  (clipped for visual convenience) is filled in with geometry from the template surface (right). (Fourth row) Underlying ground-truth surfaces of  $S_0$  (left) and  $S_2$ .

of our algorithm to missing data at the top of the head. Our algorithm not only handles the strong deformation, but also respects the details such as the mental foramen.

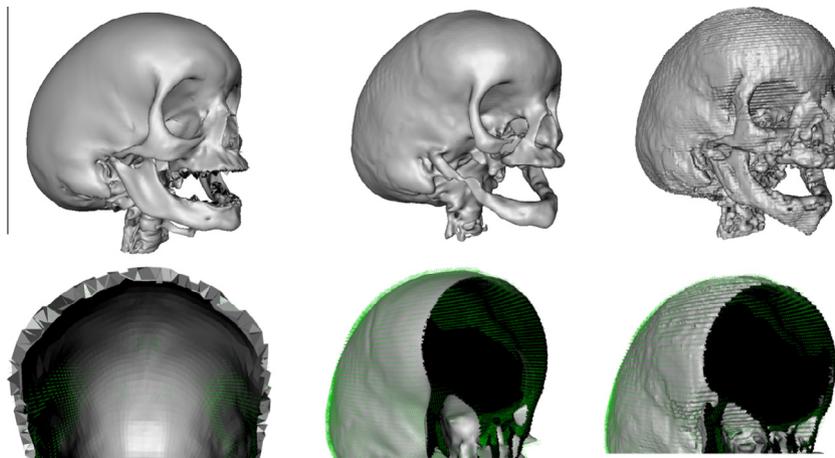
Outliers such as the CT scanner walls may occasionally appear in the data as they share similar intensity values with the skull voxels segmented from the whole image volume. Our registration performance does not degrade in such a case as exhibited for  $S_4$  in Fig. 7, where we also show another result that captures  $S_3$ . This illustrates an advantage of a template-based approach since a fully-automatic medical registration algorithm is likely to be

affected by such outliers when optimizing over the whole image volume or over the segmented binary volume (Section 7.3).

The execution times on a 2.2 GHz PC for all the experiments at high-resolution is about 20 min out of which 133 s is devoted to the initial global registration (Section 4). One deformation iteration takes about 1 s of inside vertex detection and closest point search plus 49 s of sparse Cholesky factorization, hence 50 s in total. Decimated tetrahedral mesh of 10 K vertices required 71 s of initial alignment followed by 10 s of all deformation iterations, making 81 s in total. For another decimated mesh with 5 K vertices, we



**Fig. 7.** (Top row) Initial alignments between the template mesh  $\mathcal{T}$  and  $S_3$  (two views) and  $S_4$  (rightmost). (Middle row) Resulting registered deformable model based on these initializations, one of which contains outliers ( $S_4$ ). (Bottom row) Underlying ground-truth surfaces of  $S_3$  and  $S_4$  (right).



**Fig. 8.** Target scan obtained by our method (left column), unscreened (Kazhdan and Hoppe, 2013) (middle column), and (Kazhdan and Hoppe, 2013) (right column). Necessary normal information for Kazhdan and Hoppe (2013) is transferred from our registration result (left column).

observe 25 s long initialization plus 5 s deformation, hence a total of 30 s. This reveals the efficiency of our novel volumetric deformation model. The system can easily be further sped up once the bottleneck factorization is spread on multiple cores (Hogg et al., 2010).

## 7.2. Applications

We demonstrate the potential of our algorithm by employing the registered deformable model in two important applications.

**Table 1**

Timing of our registration algorithm in comparison with medical image registration methods (Ou et al., 2011; Rueckert et al., 1999).

Method	Execution times in minutes			
	$S_1$	$S_2$	$S_3$	$S_4$
Ours	18	19	22	20
Ou et al. (2011)	617	377	548	535
Rueckert et al. (1999)	72	116	125	122

**Table 2**

Quantitative evaluation of our registration algorithm in comparison with medical image registration methods (Ou et al., 2011; Rueckert et al., 1999). Failed registrations due to wrong initial global alignment are shown as n/a. Smaller values of  $E_{reg}^{(1)}$ ,  $E_{reg}^{(2)}$ , and  $E_{reg}^{(3)}$  imply a better registration.

Method	$(E_{reg}^{(1)}, E_{reg}^{(2)}, E_{reg}^{(3)})$			
	$S_1$	$S_2$	$S_3$	$S_4$
Ours	(.21, .13, .42)	(.29, .15, .4)	(.25, .11, .38)	(.28, .1, .39)
O.11	(.22, .15, .41)	n/a	(.24, .09, .37)	(.27, .1, .36)
R.99	n/a	(.25, .11, .41)	(.22, .09, .38)	(.27, .09, .36)

### 7.2.1. Skull matching

Connectivity preservation during the volumetric registration of template  $\mathcal{T}$  gives rise to a simple yet crucial application, namely, finding correspondences between a set of ground-truth skull geometries. As discussed in Section 2, the underlying ground-truth geometry can be extracted by running an intensity-based Marching Cubes algorithm (Lorensen and Cline, 1987), which however fails to relate a given marched surface pair, say  $\mathcal{M}_j$  and  $\mathcal{M}_k$  due to  $S_j$  and  $S_k$ , respectively. To find a relationship between  $\mathcal{M}_j$  and  $\mathcal{M}_k$ , we non-rigidly register  $\mathcal{T}$  to  $S_j$  and  $S_k$  with our volumetric method, hence with better bone thickness calculation and more accurate deformation calculation. Then the closest points between the registered  $\mathcal{T}$  and  $\mathcal{M}_j$  and the registered  $\mathcal{T}$  and  $\mathcal{M}_k$  imply a correspondence between  $\mathcal{M}_j$  and  $\mathcal{M}_k$ , which we visualize in Fig. 9. In addition to general color transfer, we also track the deformation of the anthropometric landmarks  $L$  labeled on the template. As the new locations are still anatomically correct (spheres in Fig. 9), we can infer that the matching is successful. We also quantify the quality of this tracking via

$$D_{\text{grd}}(\Phi) = \frac{1}{5} \sum_{j=0}^4 D_{\text{grd}}(\phi_j), \quad (10)$$

where

$$D_{\text{grd}}(\phi_j) = \frac{1}{|\phi_j|} \sum_{(a,b) \in \phi_j} g(\vartheta(a), b), \quad (11)$$

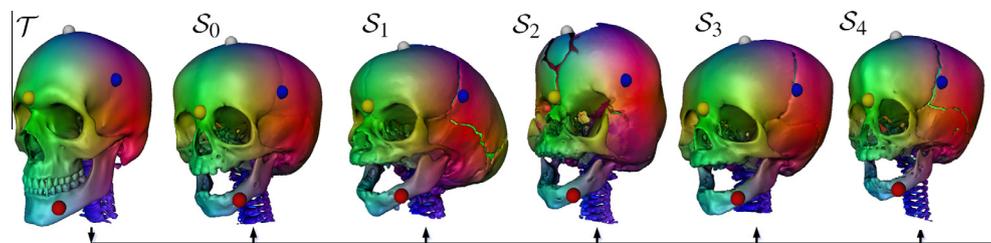
where  $(\phi_j : L \rightarrow \mathcal{M}_j) \in \Phi$  is a mapping of size  $|L|$  from the landmarks  $L$  on the template  $\mathcal{T}$  to the marching cubes surface  $\mathcal{M}_j$  of scan  $S_j$  for

$j = 0, \dots, 4$ . Here  $\vartheta(a)$  stands for the ground-truth correspondence of  $a \in L$  on the target shape  $\mathcal{M}_j$  as known a priori by manual marking, and  $g(\cdot, \cdot)$  is the geodesic distance between two points on a given surface. The maximum geodesic distance on the target model is normalized to 1 in order to simplify the interpretation of this measure. For instance, the normalized geodesic distance is 1 between the top of the head and the bottom of the neck for the surface marked with  $S_2$  in Fig. 9. Note also that one may find an insight about the results in terms of anatomical and clinical considerations by converting these numbers into millimeters (mm) once the doctor knows the geodesic distances over the actual patient. We observe  $D_{\text{grd}} = 0.033$ , which is a very small deviation from the zero error where the maximum error is 1.

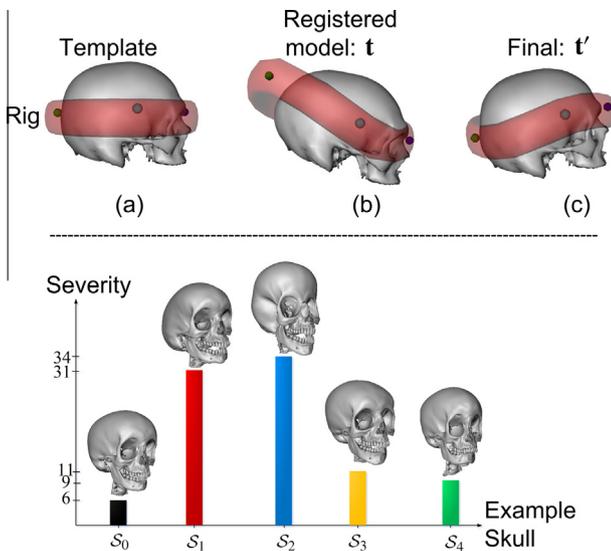
### 7.2.2. Quantification of skull deformities

There are works in the medical community that measure the severity of skull deformities by registering the pre- and post-operative CT scans of the craniosynostosis patients (Amm and Denny, 2005; Oliveira et al., 2011). Here we take a computer graphics driven approach to obtain this measure without performing any surgery. Other approaches to interpret the morphological differences rely on recognition and tracking of multiple landmarks on the skull (Lamecker et al., 2006; Marcus et al., 2008; Mendoza et al., 2013). Our approach is free of landmark matching and hence the associated problems. Our quantification process described in the sequel can also produce a plastic helmet for cranial remodeling during surgery, in case doctors decide to operate.

Rig is a simplified user interface (analogous to skeleton) that allows the user to control the deformation of the complicated surface representation (analogous to skin) easily. Vertices of the rig are bound to the related vertices of the surface mesh with blending weights such that a displacement on the simple low-resolution rig is transferred appropriately to the dense mesh, which in turn deforms the mesh in a smooth and intuitive manner. We have pre-defined this coupling between our ring-shaped low-resolution rig and high-quality template mesh (Fig. 10a) along with the blending weights. Thanks to our registration process that preserves connectivity information, this coupling between the rig and the template remains the same between the rig and the resulting registered model, and consequently user can deform the registered model of a craniosynostotic skull to the hypothetically healthy skull for quantification purposes, as shown by the two steps in Fig. 10b and c. To this effect, first the rig is positioned around the registered model  $\mathbf{t}$  accordingly (Fig. 10b). Then the interactive displacements in the rig is transferred to the anthropometric landmarks and their neighborhoods on the model. Note that there is no need for error-prone landmark matching across different skulls in this process. We use our registration result  $\mathbf{t}$  (Fig. 10b) as the rest-pose of this editing operation. The effort put in the creation of the final (hypothetically healthy) pose  $\mathbf{t}'$  (Fig. 10c) from  $\mathbf{t}$  then gives the desired severity measure, which we quantify by  $E_{\text{Dirichlet}}(\mathbf{t}')$  in Eq. (3). This interactive editing session is based on



**Fig. 9.** Transfer of colors and anthropometric landmarks  $L$  (spheres) from the template (leftmost) to all other ground-truth geometries of scans  $S_0$  to  $S_4$ . Landmarks include sinciput (yellow sphere), center of sagittal suture (gray), mental foramen (red), and center of coronal suture (blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



**Fig. 10.** Quantification of the severity of skull deformities. (Top row) Template and its transparent ring-shaped rig; note that irrelevant parts of the template are excluded for efficiency (a). Resulting registered model of our process  $\mathbf{t}$  (b) is interactively deformed to the final pose  $\mathbf{t}'$  (c). The severity score of the craniosynostotic skull  $\mathbf{t}$  is computed as  $E_{\text{Dirichlet}}(\mathbf{t}')/n$  where  $\mathbf{t}$  is used as the rest-pose and  $n$  is the number of vertices. (Bottom row) Severity scores in a bar chart. The closer this score (y-axis) to 0, the healthier the corresponding example skull (x-axis) is.

the classical as-rigid-as-possible deformation energy (Sorkine and Alexa, 2007).

**7.2.2.1. Clinical evaluation.** We have performed a minimal clinical evaluation by asking a medical doctor to use our interactive system on our five skull models. The doctor is supposed to interactively deform the registered model we provide to the hypothetically healthy skull based on her experience and the specifications of the patient, such as infant's gender, and ethnicity.

The preliminary physical examination required by the doctor is simulated by the geodesic distances over the surface of the patient. The outer surface can be extracted from the CT scan by applying the Marching Cubes algorithm (Lorenson and Cline, 1987) with

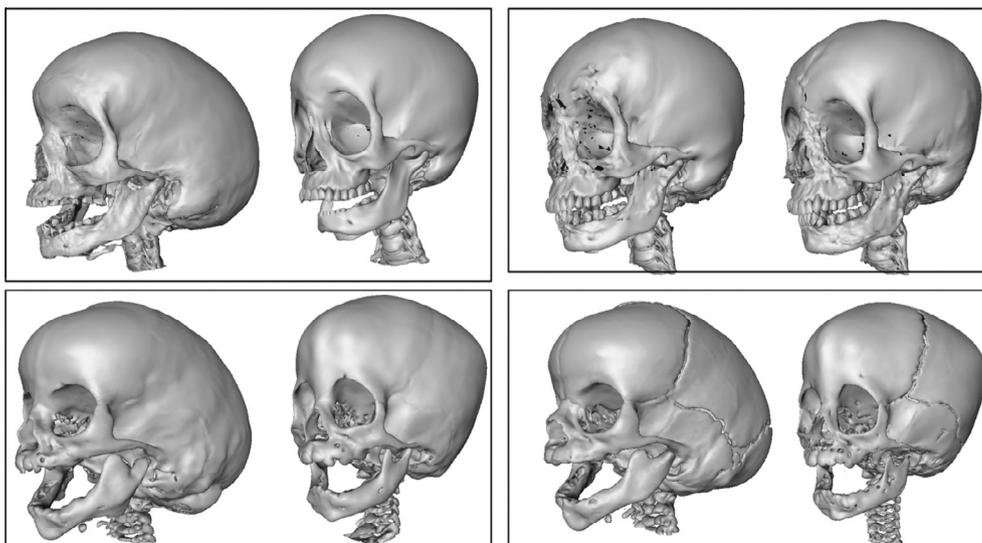
the intensity of the surface voxels (similar to the segmentation of the skull voxels which have different intensities). This simulation was necessary as we do not have physical access to the real patients.

Our doctor requested a measure for the head circumference, which should be about 36 cm for healthy infants. In order to bring the centimeter and geodesic units to the same scale, we first normalize the patient skull such that the geodesic distance between the two mental foramina along surface becomes a value prescribed by the doctor, e.g., 6.35 for  $S_0$ . We guide the interactive deformation process of the doctor by displaying the current head circumference value as well as the vertex to occiput distance to the doctor. The deformation energy used in bringing the initial skull to the hypothetically healthy skull will then be in terms of centimeters, which has a simple geometric interpretation, e.g., when divided by the number of vertices, it gives the average deformation amount of a mesh vertex. According to this evaluation, the doctor ranks the severities of our five models from high to low as  $S_2, S_1, S_3, S_4$ , and  $S_0$ , as displayed in Fig. 10 – bottom row.

Note that unlike the methods that consider only the landmarks, our quantification algorithm considers all points (Eq. (3)) and therefore does not suffer from any information loss. This severity measure is helpful in surgery decision and/or planning. Furthermore, the hypothetically healthy skull  $\mathbf{t}'$  can be 3D-printed and brought to the operating room to guide doctors while reshaping the soft skull bones of the patient.

### 7.3. Comparisons with image registration algorithms

We compare our registration result to state of the art algorithms from volumetric image registration domain (Ou et al., 2011; Rueckert et al., 1999). Rueckert et al. (1999) is a classic registration method based on image intensities and a B-spline free-form-deformation (FFD) transformation model, and still regarded as one of the best algorithms (Pszczolkowski et al., 2012). Ou et al. (2011) uses an FFD transformation model as well, but replaces the intensity-based similarity with a more sophisticated attribute-based similarity which brings higher accuracy and robustness to many tasks as both approaches are general-purpose. They, as volumetric techniques, compute deformation fields over a dense 3D image region in contrast to our elastic mesh model working solely on the manifold, which consequently makes them



**Fig. 11.** Resulting registrations for scans  $S_1$  (left in each box) and  $S_4$  (right in each box) using our method (top left box), (Ou et al., 2011) (bottom left box, left), (Rueckert et al., 1999) (bottom left box, right), and surface registration based on Rueckert et al. (1999) (top right box). Underlying ground-truth surfaces of  $S_1$  and  $S_4$  (bottom right box).

achieve our accuracy (Table 2 and Fig. 11) after about 9 h (Ou et al., 2011) or 2 h (Rueckert et al., 1999), using multi-resolution speed up. This is significantly higher than our 20 min demand for the same task (Table 1). These timing values render our method much more suitable than volumetric image registration alternatives for intra-operative CT scanning where decision making is informed during surgery without moving the patient, e.g., in the case of the rail-mounted configuration.

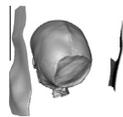
We quantify the error of the resulting registration that puts  $\mathcal{T}$  into a final orientation represented by unit normals  $\mathbf{n} = \{\mathbf{n}_1, \dots, \mathbf{n}_n\}$  as follows:

$$E_{\text{reg}}^{(1)}(\mathbf{n}, \mathbf{n}') = \frac{1}{n} \sum_{\mathbf{n}_i \in \mathbf{n}} f(\mathbf{n}_i \cdot \mathbf{n}'_i) \quad (12)$$

where vertex  $j$  of the marching cubes surface  $\mathcal{M}$  with the unit normal  $\mathbf{n}_j$  is in correspondence with the closest vertex  $i$  of the registered template  $\mathcal{T}$ , and

$$f(x) = \begin{cases} 1 - x, & x > 0 \\ 1, & x \leq 0 \end{cases} \quad (13)$$

which in turn penalizes the incompatibility of the corresponding normals between the ground-truth marched surface and the registered template, hence accounting for the thickness issue. We also provide a second measure  $E_{\text{reg}}^{(2)}$  as the average of distances between closest points of the marched and registered result where the maximum distance in average is normalized to 1. Under this normalization, Hausdorff distance between the marched and registered result constitutes our third measure  $E_{\text{reg}}^{(3)}$ . Note also that since medical image registration methods expect image data from CT scanners, we cannot feed our tetrahedral template mesh  $\mathcal{T}$  to those algorithms. Rather, we register the scan  $S_0$  to all other four scans using Ou et al. (2011) and Rueckert et al. (1999). To make the comparisons fair, we designate the source of our method as  $\mathcal{T}$  registered to  $S_0$  (Fig. 6 – second row), which effectively increases our test scope by four additional registrations. Finally note that volumetric deformation fields computed by Ou et al. (2011) and Rueckert et al. (1999) are marched (Lorensen and Cline, 1987) into surface representations for Eq. (12), Table 2 and Fig. 11. As a fourth and final error metric, we compare the volume of the skull bones in the ground-truth shape ( $V_g$ ) and in our deformed template ( $V_t$ ). Similar volumes reveal again our success in bone thickness capture (Table 3).



Running (Ou et al., 2011; Rueckert et al., 1999) on segmented binary images halves their computation times in (Table 1). The algorithms, however, become very sensitive to outliers as shown above. Scanner walls that share the same bone intensity values,

for instance, significantly degrades the registration of  $S_0$  to  $S_1$  in the back of the skull. As noted before, our template-based approach is free of such distractions (Fig. 7).

We also compare our method with a non-rigid surface registration method that models transformations using an FFD in the same way as Rueckert et al. (1999). To this end, we use as a source the boundary surface of  $\mathcal{T}$  registered to  $S_0$  (Fig. 6 – second row), and our targets become the marched surfaces of the four other patients. This simple method is fast but does not quite reach our accuracy as displayed in Fig. 11 – third row due to regularization and control point spacing issues. Note that all comparisons use the original software, kindly provided by their authors.

## 8. Limitations

This paper represents just a first step towards registration of explicitly volumetric shapes, and there are many limitations that suggest directions for future work. Although our weighting scheme that maintains the smoothness of the deforming mesh is already adaptive to capture the details, one may employ a locally-adaptive scheme in order to avoid relaxing the mesh more than is necessary. Artifacts such as inverted tetrahedra may happen occasionally. To guarantee their prevention one requires a more sophisticated energy term that penalizes nonpositive volumes (Schuller et al., 2013). This choice, however, would yield a slow nonlinear system to solve. Alternatively, regions with topological problems can be clustered to solve a local linear elastic system. Also, as we note in Fig. 4 – bottom right the volume of the tetrahedral mesh does not expand to cover the whole volumetric region it is in due to the vertices being already happy with  $E_{\text{corr}}$ . This matter of choice, although not a limitation, can be altered. Remeshing can also be incorporated so as to remove mesh parts that do not exist in the scans, such as teeth, as well as to capture the finest details, such as the cracks between the bone plates. Note that these parts to remesh can be detected automatically as the Euclidean distance between the registered template vertex and the matching closest skull voxel would yield a low confidence score around these regions. Finally, the global rigid alignment based on PCA axes is not a principled solution to the underlying abstract problem and practical performance may degrade in different scenarios.

## 9. Conclusion

This paper presents a fast and robust solution to a new version of the well-known shape registration problem, namely the registration of two volumetric shapes. The source is an explicit volumetric tetrahedral mesh, and the target is again an explicit volumetric shape represented by voxels from computed tomography (CT) scans. We adapt volumetric shape registration techniques from computer graphics such as  $E_{\text{Dirichlet}}$  to the medical imaging domain along with novel concepts added to the system such as disabling  $E_{\text{corr}}$  and adaptive weighting. Our fully-automatic volumetric shape registration algorithm has successfully modeled large non-rigid deformities of actual patients with high accuracy thanks to both the high-quality template mesh involved and consideration of volumetric issues such as skull bone thickness.

The main contribution of this research is to introduce this problem and solve it with a deformation model that respects as much volumetric information as possible. The main advantages of registering tetrahedral template over registering images are faster execution time (Table 1) and robustness to outliers (Fig. 7 – right). The side contributions are multi-initialization (Fig. 2 – left) and adaptive uniform scaling to ICP (Fig. 2 – right), matching ground-truth surfaces embedded in CT scans (Fig. 9), and objective quantification of the given medical condition craniosynostosis (Fig. 10).

**Table 3**

Proportions of the volumes of the skulls generated by the ground-truth Marching Cubes algorithm and our registration method.

	$S_0$	$S_1$	$S_2$	$S_3$	$S_4$
$V_g/V_t$	.88	.79	.81	.91	.93

An interesting future topic is to identify the most efficient user interfaces for the process of severity quantification, which may call for an example-based volumetric shape interpolation technique. Data-driven statistical models can also be investigated for the quantification purposes. Yet another interesting research direction is to develop efficient remeshing algorithms to handle the missing data, e.g., teeth, or the finest details, e.g., cracks, in the CT scans.

## Acknowledgements

We thank Jesse Goldstein for introducing us to the problems of deformed skull shapes due to craniosynostosis. We also thank Alec Jacobson for help with tetrahedral meshing and Y. Ou, A. Sotiras, and S. Psczolkowski for running the original softwares of the algorithms [Ou et al. \(2011\)](#) and [Rueckert et al. \(1999\)](#), respectively. This research was supported by TUBITAK 2219 Award and NSF CAREER Award IIS-1350330.

## Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.media.2015.03.005>.

## References

- Aiger, D., Mitra, N., Cohen-Or, D., 2008. 4-points congruent sets for robust pairwise surface registration. *Proc. SIGGRAPH*, 27.
- Allen, B., Curless, B., Popovic, Z., 2003. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Trans. Graph.* 22, 587–594.
- Amenta, N., Choi, S., Kolluri, R., 2001. The power crust. In: 6th ACM Symposium on Solid Modeling, pp. 249–260.
- Amm, C., Denny, A., 2005. Correction of sagittal synostosis using foreshortening and lateral expansion of the cranium activated by gravity: surgical technique and postoperative evolution. *Plast. Reconstr. Surg.* 116, 723–735.
- Angelov, D., Srinivasan, P., Pang, H., Koller, D., Thrun, S., Davis, J., 2004. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. *Proc. NIPS* 37, 33–40.
- Beg, M., Miller, M., Trounev, A., Younes, L., 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vision* 61, 139–157.
- Besl, P.J., McKay, N.D., 1992. A method for registration of 3D shapes. *IEEE Trans. PAMI* 14, 239–256.
- Bucki, M., Lobos, C., Payan, Y., 2010. A fast and robust patient specific finite element mesh registration technique: application to 60 clinical cases. *Med. Image Anal.* 14, 303–317.
- Cash, D., Miga, M., Sinha, T., Galloway, R., Chapman, W., 2005. Compensating for intraoperative soft-tissue deformations using incomplete surface data and finite elements. *IEEE Trans. Med. Imaging* 24, 1479–1491.
- Chang, W., Zwicker, M., 2009. Range scan registration using reduced deformable models. *Comput. Graph. Forum* 8.
- Chen, Y., Medioni, G., 1991. Object modeling by registration of multiple range images. In: *Proc. Conf. Robotics and Automation*.
- Chintalapani, G., Ellingsen, L., Sadowsky, O., Prince, J., Taylor, R., 2007. Statistical atlases of bone anatomy: construction, iterative improvement and validation. *Proc. MICCAI*, 499–506.
- Christensen, G., Joshi, S., Miller, M., 1997. Volumetric transformation of brain anatomy. *IEEE Trans. Med. Imaging* 16, 864–877.
- Chui, H., Rangarajan, A., 2003. A new point matching algorithm for non-rigid registration. *Comput. Vis. Image Und. (CVIU)* 89, 114–141.
- D'Agostino, E., Maes, F., Vandermeulen, D., Suetens, P., 2004. Non-rigid atlas-to-image registration by minimization of class-conditional image entropy. *Proc. MICCAI*, 745–753.
- Dey, T.K., Li, K., Luo, C., Ranjan, P., Safa, I., Wang, Y., 2010. Persistent heat signature for pose-oblivious matching of incomplete models. *Comput. Graph. Forum (Proc. SGP)* 29, 1545–1554.
- Ellingsen, L., Chintalapani, G., Taylor, R., Prince, J., 2010. Robust deformable image registration using prior shape information for atlas to patient registration. *Comput. Med. Imaging Graph.* 34, 79–90.
- Feldmar, J., Malandain, G., Declercq, J., Ayache, N., 1996. Extension of the icp algorithm to non-rigid intensity-based registration of 3d volumes. In: *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 84–93.
- Ferrant, M., Nabavi, A., Macq, B., Jolesz, F., Kikinis, R., Warfield, S., 2001. Registration of 3d intraoperative mr images of the brain using a finite-element biomechanical model. *IEEE Trans. Med. Imaging* 20, 1384–1397.
- Friskens, P., Rockwood, J., 2000. Adaptively sampled distance fields: a general representation of shape for computer graphics. *Proc. SIGGRAPH*, 249–254.
- Gee, J., Haynor, D., Reivich, M., Bajcsy, R., 1994. Finite element approach to warping of brain images. *SPIE Med. Imag.*, 327–337.
- Gelfand, N., Mitra, N., Guibas, L., Pottman, H., 2005. Robust global registration. *Comput. Graph. Forum (Proc. SGP)*.
- Hahnel, D., Thrun, S., Burgard, W., 2003. An extension of the icp algorithm for modeling nonrigid objects with mobile robots. *Proc. IJCAI*.
- Hogg, J., Reid, J., Scott, J.A., 2010. Design of a multicore sparse cholesky factorization using dags. *SIAM J. Sci. Comput.* 32, 3627–3649.
- Holden, M., 2008. A review of geometric transformations for nonrigid body registration. *IEEE Trans. Med. Imaging* 27, 111–128.
- Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W., 1992. Surface reconstruction from unorganized points. *Proc. SIGGRAPH*, 71–78.
- Huang, Q., Flory, S., Gelfand, N., Hofer, M., Pottmann, H., 2006. Reassembling fractured objects by geometric matching. *ACM Trans. Graph.* 25, 569–578.
- Huang, Q., Adams, B., Wicke, M., Guibas, L., 2008. Non-rigid registration under isometric deformations. *Comput. Graph. Forum (Proc. SGP)*, 1149–1158.
- Jacobson, A., Kavan, L., Sorkine, O., 2013. Robust inside-outside segmentation using generalized winding numbers. *Proc. SIGGRAPH* 32 (4), 33:1–33:12.
- Kahler, K., Haber, J., Yamauchi, H., Seidel, H.P., 2002. Head shop: Generating animated head models with anatomical structure. *Proc. SIGGRAPH*, 55–64.
- Kazhdan, M., 2007. An approximate and efficient method for optimal rotation alignment of 3d models. *IEEE Trans. PAMI*, 29.
- Kazhdan, M., Hoppe, H., 2013. Screened poisson surface reconstruction. *ACM Trans. Graph.*
- Kobbelt, L., Botshc, M., Schwanecke, U., Seidel, H.P., 2001. Feature-sensitive surface extraction from volume data. *Proc. SIGGRAPH*, 57–66.
- Lamecker, H., Zachow, S., Hege, H., Zckler, M., Haberl, H., 2006. Surgical treatment of craniosynostosis based on a statistical 3d-shape model: first clinical application. *Int. J. Comput. Assist. Radiol. Surg.*
- Li, H., Vouga, E., Gudym, A., Luo, L., Barron, J., Gusev, G., 2013. 3d self-portraits. *Proc. SIGGRAPH Asia*, 32.
- Lorensen, W., Cline, H., 1987. Marching cubes: a high resolution 3D surface construction algorithm. *Proc. SIGGRAPH*, 163–169.
- Marcus, J., Domeshek, L., Das, R., Marshall, S., Nightingale, R., Stokes, T., Mukundan, S., 2008. Objective three-dimensional analysis of cranial morphology. *Eplasty*.
- Markelj, P., Tomazevic, D., Likar, B., Pernus, F., 2012. A review of 3d/2d registration methods for image-guided interventions. *Med. Image Anal.* 16, 642–661.
- McInerney, T., Terzopoulos, D., 1995. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4d image analysis. *Comput. Med. Imaging Graph.* 19, 69–83.
- McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. *Med. Image Anal.* 1, 91–108.
- Mendoza, C., Safdar, N., Myers, E., Kittisarapong, T., Rogers, G., Linguraru, M., 2013. Computer-based quantitative assessment of skull morphology for craniosynostosis. *Lect. Notes Comput. Sci.* 7761, 98–105.
- Miller, J., Breen, D., Lorensen, W., O'Bara, R., Wozny, M., 1991. Geometrically deformed models: a method for extracting closed geometric models from volume data. *Proc. SIGGRAPH* 25, 217–226.
- Oliveira, M., Hallila, H., Ritvanen, A., Buchler, P., Paulasto, M., Hukki, J., 2011. Feature-invariant image registration method for quantification of surgical outcomes in patients with craniosynostosis: a preliminary study. *J. Pediatr. Surg.* 46, E1–E8.
- Ou, Y., Sotiras, A., Paragios, N., Davatzikos, C., 2011. Dramms: deformable registration via attribute matching and mutual-saliency weighting. *IEEE Trans. Med. Imaging* 15, 622–639.
- Pighin, F., Hecker, J., Lischinski, D., Salesin, D., Szeliski, R., 1998. Synthesizing realistic facial expressions from photographs. *Proc. SIGGRAPH*, 75–84.
- Psczolkowski, S., Pizarro, L., Guerrero, R., Rueckert, D., 2012. Nonrigid free-form registration using landmark-based statistical deformation models. *Proc. SPIE Med. Imag.: Image Process.*, 8314.
- Rohlfing, T., Maurer, C., Bluemke, D., Jacobs, M., 2003. Volume-preserving nonrigid registration of mr breast images using free-form deformation with an incompressibility constraint. *IEEE Trans. Med. Imaging* 22, 730–741.
- Rueckert, D., Sonoda, L., Hayes, C., 1999. Nonrigid registration using free-form deformations: application to breast mr images. *IEEE Trans. Med. Imaging* 18, 712–721.
- Rusinkiewicz, S., Levoy, M., 2001. Efficient variants of the ICP algorithm. *3D Digital Imaging and Modeling (3DIM)*.
- Sahillioglu, Y., Yemez, Y., 2010. Coarse-to-fine surface reconstruction from silhouettes and range data using mesh deformation. *Comput. Vis. Image Und.* 114, 334–348.
- Sahillioglu, Y., Yemez, Y., 2012. Minimum-distortion isometric shape correspondence using EM algorithm. *IEEE Trans. PAMI* 34, 2203–2215.
- Sahillioglu, Y., Yemez, Y., 2013. Coarse-to-fine isometric shape correspondence by tracking symmetric flips. *Comput. Graph. Forum* 32, 177–189.
- Schnabel, J., Tanner, C., Castellano-Smith, A., Degehard, A., Leach, M., Hose, D., Hill, D., Hawkes, D., 2003. Validation of nonrigid image registration using finite-element methods: application to breast mr images. *IEEE Trans. Med. Imaging* 22, 238–247.
- Schuller, C., Kavan, L., Panozzo, D., Sorkine-Hornung, O., 2013. Locally injective mappings. *Comput. Graph. Forum (Proc. SGP)*, 32.
- Sorkine, O., Alexa, M., 2007. As-rigid-as-possible surface modeling. *Comput. Graph. Forum (Proc. SGP)*, 109–116.
- Sotiras, A., Davatzikos, C., Paragios, N., 2013. Deformable medical image registration: a survey. *IEEE Trans. Med. Imaging* 32, 1153–1190.

- Sumner, R., Schmid, J., Pauly, M., 2007. Embedded deformation for shape manipulation. *Proc. SIGGRAPH*.
- Szeliski, R., Lavallee, S., 1996. Matching 3-d anatomical surfaces with non-rigid deformations using octree-splines. *Int. J. Comput. Vis.* 18, 171–186.
- Tagliassachi, A., Olson, M., Zhang, H., Hamarneh, G., Cohen-Or, D., 2011. Vase: volume-aware surface evolution for surface reconstruction from incomplete point clouds. *Comput. Graph. Forum (Proc. SGP)* 30, 1563–1571.
- Tam, G., Cheng, Z., Lai, Y., Langbein, F., Liu, Y., Marshall, D., Martin, R., Sun, X., Rosin, P., 2013. Registration of 3d point clouds and meshes: a survey from rigid to non-rigid. *IEEE Trans. Vis. Comput. Graph.* 19, 1199–1217.
- Taubin, G., 1995. A signal processing approach to fair surface design. *Proc. SIGGRAPH*, 315–358.
- Tevs, A., Bokeloh, M., Wand, M., Schilling, A., Seidel, H.P., 2009. Isometric registration of ambiguous and partial data. In: *Proc. Computer Vision and Pattern Recognition (CVPR)*.
- Timoner, S., 2003. Compact Representations for Fast Nonrigid Registration of Medical Images. PhD Thesis, MIT.
- Yao, J., Taylor, R., 2000. Tetrahedral mesh modeling of density data for anatomical atlases and intensity-based registration. *Proc. Med. Image Comput. Comput.-Assist. Intervent.*, 531–540.
- Zhang, H., Sheffer, A., Cohen-Or, D., Zhou, Q., van Kaick, O., Tagliasacchi, A., 2008. Deformation-driven shape correspondence. *Comp. Graph. Forum*, 1431–1439.
- Zygot, 2013. 3D Human Anatomy for Animation, Illustration, Cad and Software Development. <[www.zygot.com](http://www.zygot.com)>.