

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>

Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

A fast and robust patient specific Finite Element mesh registration technique: Application to 60 clinical cases

Marek Bucki^{a,*}, Claudio Lobos^b, Yohan Payan^{a,c}^a TIMC-IMAG Laboratory, UMR CNRS 5525, University Joseph Fourier, 38706 La Tronche, France^b Departamento de Informática, Universidad Técnica Federico Santa María, Av. Vicuña Mackenna 3939, Zip 8940897, San Joaquín, Santiago, Chile^c PIMS-Europe, UMI CNRS 3069, 200-1933 West Mall, University of British Columbia, Vancouver, BC V6T 1Z2, Canada

ARTICLE INFO

Article history:

Received 4 June 2009

Received in revised form 3 February 2010

Accepted 4 February 2010

Available online 14 February 2010

Keywords:

Mesh generation

Finite Element Method

Elastic registration

Mesh repair

ABSTRACT

Finite Element mesh generation remains an important issue for patient specific biomechanical modeling. While some techniques make automatic mesh generation possible, in most cases, manual mesh generation is preferred for better control over the sub-domain representation, element type, layout and refinement that it provides. Yet, this option is time consuming and not suited for intraoperative situations where model generation and computation time is critical. To overcome this problem we propose a fast and automatic mesh generation technique based on the elastic registration of a generic mesh to the specific target organ in conjunction with element regularity and quality correction. This Mesh-Match-and-Repair (MMRep) approach combines control over the mesh structure along with fast and robust meshing capabilities, even in situations where only partial organ geometry is available. The technique was successfully tested on a database of 5 pre-operatively acquired complete femora CT scans, 5 femoral heads partially digitized at intraoperative stage, and 50 CT volumes of patients' heads. In the latter case, both skin and bone surfaces were taken into account by the mesh registration process in order to model the face muscles and fat layers. The MMRep algorithm succeeded in all 60 cases, yielding for each patient a hex-dominant, Atlas based, Finite Element mesh with submillimetric surface representation accuracy, directly exploitable within a commercial FE software.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Physically based models are now widely used in the field of biomedical engineering to represent human organs' geometrical and mechanical behaviors. Among them, numerical models based on the Finite Element Method (Hughes, 1987) became very popular because of their ability to address the complex geometries, the anisotropic material properties and the specific boundary conditions associated with living tissues.

In the field of medical imaging, the first Finite Element (FE) models were mostly used to better understand and validate a given surgical treatment, to model physiological behaviors or to provide virtual simulators for clinicians. In these frameworks, models were limited to a single generic model for each study (the terminology "Atlas" was often used).

More recently, applications in the domain of Computer Assisted Planning and Computer Aided Surgery sparked the need for patient-specific FE models representing the modeled organ geometry reconstructed from patient medical image data, such as computed

tomography (CT) or magnetic resonance imaging (MRI). In most cases, organs are identified in these data sets by means of manual, semi-automatic or automatic segmentation tools that extract shape information (3D points, contours and/or surfaces) necessary for the generation of the Finite Element mesh representing the volume of the organ.

For both segmentation and FE mesh generation phases, manual intervention is often required which can make this procedure long and tedious. This is especially true in situations where the pre- or intraoperative time window or clinician availability to perform these delicate tasks is limited.

This paper addresses the second phase, namely the FE mesh generation, with the introduction of a procedure – the Mesh-Match-and-Repair (MMRep) algorithm – that allows a fast and fully-automatic patient-specific FE mesh generation.

Although FE models were already used in the field of biomedical engineering at the beginning of the 1980s (Jaspers et al., 1980; Huiskes, 1980), researchers only began focusing on automatic mesh generation in the early 2000s (Couteau et al., 2000; Gibson et al., 2003; Viceconti et al., 2004; Taddei et al., 2004; Luboz et al., 2005; Liao et al., 2005; Shim et al., 2007; Sigal et al., 2008; Grosland et al., 2009). The primary motivation for automatic mesh generation algorithms was that patient specific FE models could be

* Corresponding author.

E-mail address: Marek.Bucki@imag.fr (M. Bucki).

routinely used by the clinicians. In that domain, the vast majority of automatic mesh generators for living tissues produce tetrahedral meshes (Molino et al., 2003; George et al., 2002; Alliez et al., 2005; Si and Gaertner, 2005). Based on a 3D surface representing the external geometry of the organ, these algorithms produce a volume made of high quality 3D tetrahedra and sometimes allow for adaptive mesh refinement (Si, 2006, Tetgen¹).

Some researchers in biomechanics continue to propose hand-made FE meshes (Chabanas and Payan, 2000; Luboz et al., 2004; Wittek et al., 2007), mainly for two reasons. Firstly, they argue that it is important to be able to identify sub-regions associated to anatomical sub-structures inside the 3D FE mesh (e.g. the ventricles, the tumor and the hemispheres of the brain FE mesh proposed by Wittek et al. (2007)). These sub-regions, corresponding to sets of elements, are labeled and associated with specific constitutive behaviors and boundary conditions. Secondly, they tend to prefer hexahedra over tetrahedra, based on numerical considerations (Benzley et al., 1995) as well as the fact that for incompressible and/or nearly incompressible materials, 4-noded tetrahedra with linear shape functions tend to lock and become overly stiff (Hughes, 1987). Despite some improvements to this method of creating FE meshes this work still requires an excessive amount of manual effort to achieve satisfactory results.

In order to still benefit from this manual design while providing automatic FE mesh generation, techniques termed “registration methods” or “morphing methods” were recently proposed (Couteau et al., 2000; Castellano-Smith et al., 2001; Sigal et al., 2008; Grosland et al., 2009). The main premise is to start with a predefined “generic” (or “template”) FE mesh that represents the organ. This mesh is manually designed to include any necessary sub-regions and hexahedra, and to preserve element quality, orientation and density in regions that require it for numerical simulation. This template is then automatically morphed onto the patient data (3D landmark points, contours and/or surfaces) that was extracted from the segmentation of the medical images. This process generates a patient-specific mesh adapted to the geometries of the anatomical structures extracted from patient data, with a mesh topology that is similar to that of the template (same nodes and elements organization).

Our group was the first to initiate this principle of mesh morphing with the introduction of the Mesh-Matching algorithm (Couteau et al., 2000). Since then, we encountered two strong limitations with our mesh morphing principle. Both are due to the fact that the template mesh quality after registration can be strongly decreased when the morphing algorithm induces excessive spatial distortions. Therefore, the first consequence is the inability to maintain the regularity of some elements, which disables FE analysis from being carried out. This concern was partially discussed in Luboz et al. (2005). The second consequence of the morphing method is that the elements shape qualities are decreased in some regions of the template mesh, which leads to lower accuracy in the numerical simulation (Field, 2000; Kwok and Chen, 2000; Shewchuk, 2002).

This paper aims at introducing our latest algorithms concerning (1) the elastic registration method that guarantees a C^1 -diffeomorphic transform; and (2) the mesh repair technique that ensures that the produced mesh complies with both regularity and quality criteria. The MMRep approach was applied and evaluated on 60 clinical cases, which is, to our knowledge, the largest database ever tested in the literature for a patient-specific FE mesh registration method.

The MMRep algorithm is a two-step sequential procedure. Firstly, the patient data and the Atlas mesh surface nodes are

registered using the **elastic deformation procedure** described in Section 2. This deformation is then applied to the inner Atlas nodes yielding a FE mesh that represents the modeled domain with sufficient accuracy. As a consequence of this deformation, the Atlas elements may suffer distortions and become either “irregular elements” which make FE analysis impossible, or “poor quality elements” in which case the computation, although feasible, can exhibit numerical instabilities. To recover mesh regularity and reach an acceptable quality level, the **mesh repair procedure** described in Section 3 is carried out on the deformed mesh.

2. Elastic registration

2.1. Registration overview

We define an elastic registration function as a mapping $R: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that superimposes a source point cloud S onto a target, or “destination”, data set D , which can either be a point cloud or a surface mesh. The computed elastic registration procedure complies with continuum mechanics conditions on motion (Belytschko et al., 2006) as R defines a C^1 -**diffeomorphic, non-folding and one-to-one** correspondence between geometric data sets, as demonstrated, respectively, in Sections 2.2–2.4.

The input source points set is initially embedded in a deformable “virtual elastic grid”. We arbitrarily set the shape of the grid to be the bounding box of the points, extended by a 10% margin. The considered deformation R is formed by successive elementary grid deformations noted r , all having the desired regularity properties, much in the same way as proposed in Rueckert et al. (2006).

The regular grid is progressively refined in order to increase registration accuracy and the expression of the compound registration function is:

$$R = \underbrace{r_j^{N_j} \circ \dots \circ r_j^1}_{G_j} \circ \dots \circ \underbrace{r_1^{N_1} \circ \dots \circ r_1^1}_{G_1}$$

where G_j , $j = 1, \dots, J$ are successive grid refinements, and r_j^i , $i = 1, \dots, N_j$, elementary deformations performed at grid level j . As level indications are irrelevant in the following demonstrations, the deformation R is simply written as:

$$R = r_N \circ \dots \circ r_1 \tag{1}$$

where $N = \sum_{j=1}^J N_j$.

If required, the inverse registration, R^{-1} , can be computed by combining the elementary inverses in the reverse order of the direct registration, thus: $R^{-1} = r_1^{-1} \circ \dots \circ r_N^{-1}$.

At each step n , the choice of the elementary deformation r_n to be applied to the source data is driven by the minimization of a “registration energy” E which measures the similarity between the deformed source points set and the destination data set D . As geometrical shape similarity is sought, E is defined as the sum of Euclidean distances between S and D :

$$E(r_n \circ R_{n-1}) = \sum_{\mathbf{s} \in S} d(r_n(R_{n-1}(\mathbf{s})), D) \tag{2}$$

where $R_{n-1} = r_{n-1} \circ \dots \circ r_1$ represents the registration function assembled at step $n - 1$, and the operator $d(\cdot, \cdot)$ can either be a point-to-point or a point-to-surface distance measure, depending on the nature of D .

In order to speed up the computations of energy E , a distance map is generated from the destination data set D prior to registration (Saito and Toriwaki, 1994). Distance map voxel dimensions are $1 \times 1 \times 1$ mm in order to achieve submillimetric surface repre-

¹ <http://tetgen.berlios.de>.

sentation. The sampled space region covers the bounding box of the considered destination data, extended by a 5% margin.

Point-to-surface registrations are performed in all three use cases discussed in Section 4. For each destination Atlas mesh, signed point-to-surface distances are measured using surface orientation information. A positive distance is recorded in the distance map for points lying outside the closed destination surface, and a negative distance is recorded for points lying inside. The distance map computations can take several minutes, which is not a limitation to our approach, even in the intraoperative context illustrated in Section 4.2, as the Atlas meshes are processed pre-operatively. At registration time, source-to-destination distance evaluations are done by trilinear interpolation within the signed distance map and the absolute value of the result is retained.

The r_n functions are successively chosen so that E decreases optimally at each step. To this end the virtual grid is subdivided into a number of regular hexahedrons called “cells”. Each node of the grid is considered separately and the gradient of the registration energy is computed as function of the node’s position. The opposite of this vector defines the node’s preferred displacement, that is, the one that will lead to the greatest registration energy decrease achievable by moving the considered node.

Let $S_0 := S$ be the initial source points set, and let S_i be the source points set at iteration i . Of all nodes, the one leading to the highest energy decrease is chosen and its preferred displacement is applied while all the other nodes remain fixed. This nodal displacement is propagated throughout the neighboring grid cells and the affected source points are moved accordingly to generate S_{i+1} , the new set of deformed source points.

Once this basis deformation step applied, the virtual grid returns to its initial regular configuration and the source points set S_{i+1} is embedded within it. A new iteration can be computed by taking the new configuration S_{i+1} as input. This Eulerian strategy allows large deformations of S to be achieved without having to maintain large-strain consistency of the virtual mesh, as would be the case if the virtual grid strictly followed the deformation with each iteration.

The regularity of the grid before each iteration also saves computational time by allowing the interpolation of a node displacement throughout its neighboring cells to be computed using a template unitary cell. The iterations stop when no significant energy decrease can be achieved by moving any grid nodes at the current grid refinement level. In our implementation, we have set this stop threshold T_E to 1%, and the registration iterations continue as long as energy E can be reduced by more than 1%, i.e. $(E_i - E_{i+1})/E_i > T_E$. Thus, the number of iterations performed at each refinement step is variable.

The above iterative loop describes the procedure at a given grid refinement level. In order to maintain the spatial consistency of S during the assembly of the registration transformation, a top-down hierarchical approach has been implemented. The iterative assembly of the registration function R starts at the coarsest grid level.

Once the deformation search at the current level has been exhausted, the grid is refined by subdividing each cell into eight smaller ones in an octree method.

Fig. 1 illustrates in 2D the multi-grid iterative registration technique. The source points S are represented by the grey dots; for clarity D is not shown. In (a) the initial set $S_0 := S$ is embedded within the virtual square grid discretized at the coarsest level 1, and the energy gradients are computed at the 4 mesh nodes. In (b) the optimal node displacement is applied to the grid and the source points move producing a new source set S_1 . If no significant energy decrease can be generated at level 1, the grid is refined at level 2, (c), S_1 is embedded within it, and the energy gradients are computed at the 9 nodes of the mesh. In (d) the best nodal displacement is applied and the resulting source points set S_2 is computed.

At each grid refinement level, the size of the registered source features is approximately that of the current grid cell. The virtual grid nodal displacement is only applied if the resulting grid deformation leads to a registration energy decrease. As all the source points located in the grid cells surrounding the displaced node are altered, the source features significantly smaller than the current cell size have limited impact on the registration sequence. The dimensions of the smallest cell reached during the top-down refinement descent thus roughly define the size of the specific features present in S that may not be registered if the corresponding feature is absent from D .

By primarily focusing on larger source features, this hierarchical approach also reduces the influence of noise usually found in patient data gathered from pre-operative medical image segmentation (see Section 4.1) or from intraoperative acquisitions (see Section 4.2). During the evaluation of the MMRep technique the smallest cell size was set to 1 mm, which required about 8 or 9 grid subdivision levels as the typical model size was about 25 cm and $250 \text{ mm}/2^8 \approx 1 \text{ mm}$.

Furthermore, the mechanical regularization strategy described in Section 2.6 limits excessive space distortions due to the presence of noise by monitoring the magnitude of potential elastic energy in the source space throughout the deformation process. Accurate registration of outliers has indeed a prohibitive mechanical cost which, in our procedure, discards elementary deformations attempting to register them.

The optimization procedure is done by a gradient descent technique (Press et al., 1992). The energy gradients are evaluated at each node by the Finite Differences method. The line search in the direction of the gradient descent is performed by approximating the energy curve by a parabola $P(t) = at^2 + bt + c$. The parameters a , b and c are deduced from: the current value of E at the considered node, corresponding to $P(0)$; the slope of E in the descent direction, defining $P'(0)$; and the value of E at the maximally displaced node position (see Section 2.3), yielding $P(1)$. Finally, the optimal descent step in the current direction is given by $t_{\min} = -b/(2a)$.

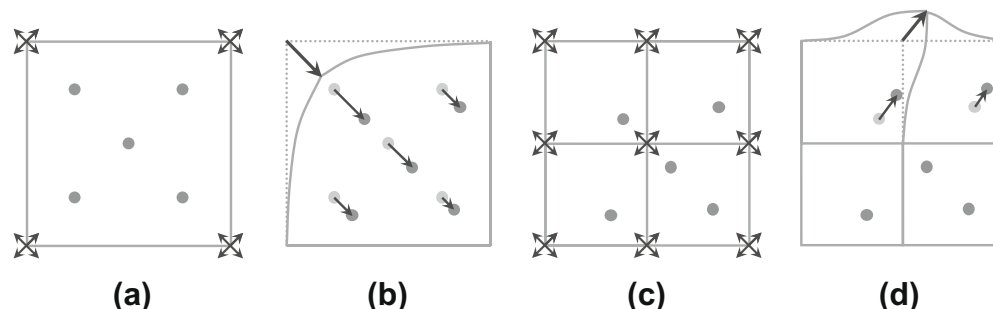


Fig. 1. Elastic registration overview. (a) S_0 at refinement level 1. (b) S_1 after deformation at level 1. (c) S_1 at refinement level 2. (d) S_2 after deformation at level 2.

2.2. C^1 -differentiability of the deformation

At each registration step, the displacement applied to a given grid node \mathbf{n} is propagated to the source points located in cells surrounding \mathbf{n} by means of a “weight” function $w_n : \mathbb{R}^3 \rightarrow [0, 1]$.

Let \vec{u} be the displacement applied to node \mathbf{n} , and \mathbf{s} be a source point found in a cell affected by the movement of \mathbf{n} . The displacement propagated to \mathbf{s} is $w_n(\mathbf{s})\vec{u}$. Similarly to the shape functions in the Finite Element Method, the support of the weight function w_n is the union of the cells neighboring \mathbf{n} . The deformation consistency is further ensured by the two following conditions: $w_n(\mathbf{n}) = 1$ and $w_n = 0$ at the boundary of its support.

From the above, it follows that the elementary elastic registration function r created by the displacement \vec{u} of node \mathbf{n} has the following expression:

$$r : \mathbb{R}^3 \rightarrow \mathbb{R}^3, \mathbf{s} \mapsto \mathbf{s} + w_n(\mathbf{s})\vec{u} \tag{3}$$

C^1 -differentiability of the elementary registration function r stems from the C^1 -differentiability of the weight function w_n and the Jacobian matrix of r is given by:

$$J_r := \frac{\partial r}{\partial \mathbf{s}} = Id_{3 \times 3} + \begin{pmatrix} u_x \\ u_y \\ u_z \end{pmatrix} \begin{pmatrix} \frac{\partial w_n}{\partial x} & \frac{\partial w_n}{\partial y} & \frac{\partial w_n}{\partial z} \end{pmatrix} \tag{4}$$

Now, let $\mathbf{s}_0 \in \mathbb{R}^3$ be an arbitrary source point and \mathbf{s}_n this point transformed after applying n elementary registration steps. Then $\{\mathbf{s}_n, n \in [0, N]\}$ is the set of all successive positions of \mathbf{s}_0 and, according to Eq. (1), $R(\mathbf{s}_0) = \mathbf{s}_N$. The Jacobian matrix of R computed at point \mathbf{s}_0 is given by the following chain rule:

$$J_R(\mathbf{s}_0) := \frac{\partial R}{\partial \mathbf{s}}(\mathbf{s}_0) = \frac{\partial r_N}{\partial \mathbf{s}}(\mathbf{s}_{N-1}) \dots \frac{\partial r_2}{\partial \mathbf{s}}(\mathbf{s}_1) \frac{\partial r_1}{\partial \mathbf{s}}(\mathbf{s}_0) \tag{5}$$

All weight functions w_n stem from a “template” weight function “ w ” defined on the $[0, 1]^3$ grid cell and associated to node $(1,1,1)$. Let π be a third degree polynomial defined by $\pi(t) = t^2(3 - 2t)$, such as $\pi(0) = 0, \pi(1) = 1, \pi'(0) = 0$ and $\pi'(1) = 0$. The template weight function w is defined on $[0, 1]^3$ as:

$$w(\mathbf{s}) = w(s_1, s_2, s_3) := \pi(s_1)\pi(s_2)\pi(s_3) \tag{6}$$

A specific weight function w_n is defined on the union of neighboring cells around the displaced node \mathbf{n} by variable change and scaling in order to adapt the canonic $[0, 1]^3$ domain to the cell size within the actual grid. Fig. 2a illustrates on a 2D grid the variable changes needed to construct from w the weight function w_n , with $\mathbf{n} = (1, 1)$, defined on a 2×2 cells neighborhood of \mathbf{n} . Fig. 2b shows the 3D plot of the resulting weight function $w_n : \mathbb{R}^2 \rightarrow [0, 1]$.

The two other regularity properties of the elementary registration functions r , namely **bijection** and **non-folding**, are enforced by limiting the amplitude of nodal displacements, as described in the following sections.

2.3. Control over space distortion

Space **folding** can be mathematically expressed as follows. Consider a positively oriented set of three infinitesimal vectors, $d\mathbf{X}_1, d\mathbf{X}_2$ and $d\mathbf{X}_3$, placed in undeformed space at point \mathbf{s} and defining the volume dV . The signed value of dV can be computed as the determinant of the matrix formed by the three vectors.

$$dV = (d\mathbf{X}_1 \times d\mathbf{X}_2) \cdot d\mathbf{X}_3 = |d\mathbf{X}_1 \ d\mathbf{X}_2 \ d\mathbf{X}_3|$$

R is said to be **locally non-folding** if the deformed infinitesimal volume $dv = R(dV)$ defined by the three deformed vectors dx_1, dx_2 and dx_3 , remains positive, or, in other words, if R does not change the orientation of space in the neighborhood of \mathbf{s} .

The differential form of R gives us the expression of the deformed vectors:

$$\forall i = 1, 2, 3, d\mathbf{x}_i = \frac{\partial R}{\partial \mathbf{s}}(\mathbf{s}) d\mathbf{X}_i$$

and the relation between dv and dV is:

$$dv = |d\mathbf{x}_1 \ d\mathbf{x}_2 \ d\mathbf{x}_3| = \left| \frac{\partial R}{\partial \mathbf{s}}(\mathbf{s}) \right| |d\mathbf{X}_1 \ d\mathbf{X}_2 \ d\mathbf{X}_3| = |J_R(\mathbf{s})| dV$$

where $J_R(\mathbf{s}) := (\partial R / \partial \mathbf{s})(\mathbf{s})$ is the Jacobian of R given by the chain rule in Eq. (5). It follows that an application R is **non-folding** if its Jacobian is strictly positive, i.e.:

$$\forall \mathbf{s} \in \mathbb{R}^3, |J_R(\mathbf{s})| > 0 \tag{7}$$

If the value of the Jacobian at a given point is greater than 1, the deformation locally **stretches** space and if its value is smaller than 1, space is locally **compressed**.

From Eq. (5) it follows that R is non-folding if and only if all $r_i, i = 1, \dots, N$ are non-folding, since:

$$|J_R| = \left| \frac{\partial R}{\partial \mathbf{s}} \right| = \prod_{i=1}^N \left| \frac{\partial r_i}{\partial \mathbf{s}} \right| = \prod_{i=1}^N |J_{r_i}|$$

Let's now see how to ensure that each individual elementary registration function is non-folding. Using Eq. (4), the non-folding condition on r becomes:

$$\forall \mathbf{s} \in \mathbb{R}^3, |J_r(\mathbf{s})| = 1 + \vec{\nabla} w_n(\mathbf{s}) \cdot \vec{u} > 0 \tag{8}$$

Given the above expression of the template weight function w , the magnitude of the gradient of a specific weight function w_n defined




Fig. 2. (a) Variable changes required in 2D for the assembly of w_n on a neighborhood of 4 cells centered on node $\mathbf{n} = (1, 1)$. (b) Value of w_n plotted over the four 2D cells. (c) An elementary deformation leaves all segments [a, b] parallel to the applied nodal displacement unchanged.

on a $[0, L]^3$ grid cell has an upper bound of $3\sqrt{3}/(2L)$. To ensure non-folding it is thus sufficient to limit the amplitude of nodal displacements:

$$\|\vec{u}\| \leq 38\% L < \frac{2}{3\sqrt{3}} L$$

In our implementation of the algorithm this value has been reduced to **10% of the current grid cell size**, so as to not only achieve non-folding but also limit space distortion at each elementary registration step. As $\|\vec{u}\| \leq L/10$ and $\|\vec{\nabla} w_n\| \leq 3\sqrt{3}/(2L)$, it follows that:

$$1 - \frac{3\sqrt{3}}{20} \leq 1 + \vec{\nabla} w_n(\mathbf{s}) \cdot \vec{u} \leq 1 + \frac{3\sqrt{3}}{20} \quad (9)$$

The above equation gives us, for all cell sizes, approximate² Jacobian lower and upper bounds: $0.74 < |J_r| < 1.26$, and hence $0.74^N < |J_R| < 1.26^N$.

2.4. Registration inversion

In this section we will show that under the non-folding constraint, the deformation R is a bijection and discuss how its inverse, R^{-1} , can be computed with a pre-defined level of accuracy for any point $\mathbf{q} \in \mathbb{R}^3$.

The registration function R is one-to-one if the same is true of each elementary registration r . To prove that a non-folding elementary registration function is also a bijection, consider a 2D deformation \vec{u} , as defined in Eq. (3), created by applying the displacement \vec{u} to the central node of the 4 cells group depicted in Fig. 2c.

Given that all the displacements applied to the source points are collinear with \vec{u} , parallel segments such as $[\mathbf{a}, \mathbf{b}]$ in Fig. 2c, map onto themselves. Furthermore, points lying outside or on the boundary of the 4-cell group are left unchanged. To prove that r is a bijection it is thus sufficient to show that it defines a one-to-one application $[\mathbf{a}, \mathbf{b}] \rightarrow [\mathbf{a}, \mathbf{b}]$, for any segment $[\mathbf{a}, \mathbf{b}]$ parallel to \vec{u} .

As the considered segment is parallel to \vec{u} , there exists a scalar β such as $\mathbf{b} = \mathbf{a} + \beta\vec{u}$. Moreover, all segment points $\mathbf{p} \in [\mathbf{a}, \mathbf{b}]$ can be written as $\mathbf{p} = \mathbf{a} + \rho\vec{u}$, with $\rho \in [0, \beta]$, leading to $r(\mathbf{p})$ being rewritten as follows:

$$r(\mathbf{p}) = \mathbf{p} + w_n(\mathbf{p})\vec{u} = \mathbf{a} + (\rho + w_n(\mathbf{a} + \rho\vec{u}))\vec{u}$$

The deformation r is thus a bijection $[\mathbf{a}, \mathbf{b}] \rightarrow [\mathbf{a}, \mathbf{b}]$ if and only if the mapping $f : \rho \mapsto \rho + w_n(\mathbf{a} + \rho\vec{u})$ is also a bijection $[0, \beta] \rightarrow [0, \beta]$, which we shall prove is true under the non-folding hypothesis.

Deformation r leaves the segment extremities unchanged, $r(\mathbf{a}) = \mathbf{a}$ and $r(\mathbf{b}) = \mathbf{b}$, which, using the above expressions, leads to $f(0) = 0$ and $f(\beta) = \beta$, respectively. Furthermore, the derivative of f is $f' = 1 + \vec{\nabla} w_n \cdot \vec{u}$. As a consequence, if r is non-folding then, according to Eq. (8), we have $f' > 0$ which implies that the strictly monotonic function f is a bijection $[0, \beta] \rightarrow [0, \beta]$ and so is $r : [\mathbf{a}, \mathbf{b}] \rightarrow [\mathbf{a}, \mathbf{b}]$, which concludes our proof.

Another issue is computing the inverse registration $R^{-1} = r_1^{-1} \circ \dots \circ r_N^{-1}$ with a user-defined accuracy ϵ , measured in undeformed space. The task consists in finding, for a given deformed point $\mathbf{q} \in \mathbb{R}^3$, an undeformed point $\mathbf{p} \in \mathbb{R}^3$, such as $\|\mathbf{p} - R^{-1}(\mathbf{q})\| < \epsilon$. We will first show how to accurately compute the inverse of an elementary registration function r , and then use these results to accurately inverse the compound registration R .

Given $\mathbf{q} \in [\mathbf{a}, \mathbf{b}]$, $\mathbf{q} = \mathbf{a} + \mu\vec{u}$, finding $\mathbf{p} = \mathbf{a} + \rho\vec{u}$ such as $\mathbf{p} = r^{-1}(\mathbf{q})$ reduces to solving $\rho = f^{-1}(\mu)$ on $[0, \beta]$. The mapping f is a 9th degree polynomial (see Eq. (6)) and as no analytical expression of the solution is available, it must be approximated iteratively using, for example, a Newton–Raphson procedure.

Let $\vec{p} = \mathbf{a} + \vec{p}\vec{u} = r^{-1}(\mathbf{q})$ be the searched point, and $\mathbf{p} = \mathbf{a} + \rho\vec{u} \approx r^{-1}(\mathbf{q})$ its iteratively computed approximation. The approximation error in undeformed space can be rewritten as:

$$\|\mathbf{p} - r^{-1}(\mathbf{q})\| = \|\mathbf{p} - \vec{p}\| = |\rho - \vec{p}|\|\vec{u}\| \quad (10)$$

In order to compute the inverse registration with the desired accuracy level ϵ , the exact solution $\vec{p} = f^{-1}(\mu)$ must be approximated by ρ so that:

$$|\rho - \vec{p}| < \frac{\epsilon}{\|\vec{u}\|}$$

As the value of \vec{p} is unknown, the approximation error must be computed in deformed space. To this end the above expression is rewritten and, using the finite increments theorem, yields:

$$|\rho - \vec{p}| = |f^{-1}(f(\rho)) - f^{-1}(\mu)| < M |f(\rho) - \mu| \quad (11)$$

The above constant M is computed using the relation $(f^{-1})' = (f')^{-1}$, which in conjunction with Eq. (9), gives $1/1.26 < (f^{-1})' < 1/0.74 = M$. This leads us to the conclusion that in order to compute an ϵ -accurate inverse of \mathbf{q} in undeformed space, Newton–Raphson iterations must be carried out until the approximate parameter ρ satisfies:

$$|f(\rho) - \mu| < 0.74 \frac{\epsilon}{\|\vec{u}\|}$$

Now let us compute the inverse of a compound registration function R . To control the accumulated error overhead at each elementary inversion step we will use the fact, demonstrated by combining Eqs. (10) and (11), that the inverse of any elementary registration function r is **M–Lipschitz** continuous, i.e.:

$$\forall \mathbf{q}, \mathbf{p} \in \mathbb{R}^3, \|r^{-1}(\mathbf{q}) - r^{-1}(\mathbf{p})\| < M \|\mathbf{q} - \mathbf{p}\|$$

For the sake of clarity, let $R = r_3 \circ r_2 \circ r_1$ be the considered registration function and $\mathbf{s}_3 \in \mathbb{R}^3$ a point in deformed space which inverse, $R^{-1}(\mathbf{s}_3)$, is sought with accuracy ϵ . Adopting the notation convention used to derive Eq. (5) above, we have:

$$R^{-1}(\mathbf{s}_3) = (r_1^{-1} \circ r_2^{-1} \circ r_3^{-1})(\mathbf{s}_3) = (r_1^{-1} \circ r_2^{-1})(\mathbf{s}_2) = r_1^{-1}(\mathbf{s}_1) = \mathbf{s}_0$$

Fig. 3 illustrates the 3-step computation of $R^{-1}(\mathbf{s}_3)$ along with error accumulation due to approximations performed at each elementary inversion step. Undeformed space is represented on the left, and deformed space on the right of the figure. Approximation steps are shown as dashed oblique arrows, and exact inversions as horizontal arrows.

We shall now establish the relation between ϵ , the selected tolerance in undeformed space, and successive elementary approximation errors ϵ_3 , ϵ_2 and ϵ_1 .

Inversion of r_3 . $\mathbf{s}_2 = r_3^{-1}(\mathbf{s}_3)$ is approximated by \mathbf{t}_2 with accuracy ϵ_3 in r_3 -undeformed space: $\epsilon_3 > \|\mathbf{t}_2 - \mathbf{s}_2\|$. Using the Lipschitz




Fig. 3. Approximation of $\mathbf{s}_0 = R^{-1}(\mathbf{s}_3)$ by \mathbf{v}_0 . Horizontal arrows represent exact inversions and oblique dashed arrows iterative inverse approximations. Left: undeformed space, right: deformed space.

² $3\sqrt{3}/20 \approx 0.259$.

constant M , this error propagates to the left as: $\epsilon_3 M > \|\mathbf{t}_1 - \mathbf{s}_1\|$, and $\epsilon_3 M^2 > \|\mathbf{t}_0 - \mathbf{s}_0\|$.

Inversion of r_2 . $\mathbf{t}_1 = r_2^{-1}(\mathbf{t}_2)$ is approximated by \mathbf{u}_1 with accuracy ϵ_2 in r_2 -undeformed space: $\epsilon_2 > \|\mathbf{u}_1 - \mathbf{t}_1\|$. As in the previous step, this error propagates to the left as: $\epsilon_2 M > \|\mathbf{u}_0 - \mathbf{t}_0\|$.

Inversion of r_1 . Finally, $\mathbf{u}_0 = r_1^{-1}(\mathbf{u}_1)$ is approximated by \mathbf{v}_0 with accuracy ϵ_1 in initial undeformed space: $\epsilon_1 > \|\mathbf{v}_0 - \mathbf{u}_0\|$.

The final error in undeformed space $\|\mathbf{v}_0 - \mathbf{s}_0\|$ can be upper bound as:

$$\|\mathbf{v}_0 - \mathbf{s}_0\| \leq \|\mathbf{v}_0 - \mathbf{u}_0\| + \|\mathbf{u}_0 - \mathbf{t}_0\| + \|\mathbf{t}_0 - \mathbf{s}_0\| < \epsilon_1 + \epsilon_2 M + \epsilon_3 M^2$$

The above expression can be generalized to any compound registration function $R = r_N \circ \dots \circ r_1$. In order to meet the desired accuracy standard ϵ , the computational effort can be spread among all N elementary inversions by setting the Newton–Raphson approximation threshold for each r_n^{-1} , $n = 1, \dots, N$ to:

$$\epsilon_n = \frac{\epsilon}{N M^{n-1}}$$

This is a reasonable heuristic as ϵ_n is smaller for higher values of n , corresponding to finer grid levels and hence, smaller nodal displacements (10% of grid cell size – see Section 2.3). Solution search intervals are thus narrower, and higher inversion accuracy is thus easier to achieve than at coarser grid levels.

There is a final consideration on registration inversion: if the cubic polynomials π used to define the shape function w in Eq. (6) are replaced with the identity $i : [0, 1] \rightarrow [0, 1], t \rightarrow t$, then f becomes a 3rd degree polynomial and the exact inverse of the registration function R can be computed in a single iteration by using the well-known Tartaglia–Cardan formula. Nevertheless, this enhancement comes at a price as the smoothness of the registration function R is degraded from C^1 to C^0 .

2.5. Asymmetric registration

The “asymmetric” energy E defined in Eq. (2) handles situations where the features present in S are present in D but the reciprocal is not necessarily true. A situation can occur in which the Atlas mesh only represents a fraction of the organ features available within the patient data (see Section 4.3), requiring that the elastic registration R is computed as transforming the Atlas to fit the patient data. The patient specific mesh is then obtained by application of this transform to the generic mesh, i.e. **Patient mesh = $\mathbf{R}(\text{Atlas})$** .

Conversely, if only partial information about the patient’s organ is available, its overall shape has to be inferred from the a priori knowledge carried by the generic Atlas mesh (see Sections 4.1 and 4.2). This is done by taking the patient data as source points, and computing the registration function R that fits the source points onto the Atlas. The patient specific mesh is obtained by inverting the resulting elastic deformation, i.e. **Patient mesh = $\mathbf{R}^{-1}(\text{Atlas})$** .

2.6. Mechanical regularization

In order to avoid excessive space distortions, a mechanical regularization term is monitored during the registration process, thereby upgrading the virtual elastic “grid” concept to virtual elastic “solid”. As the elastic registration compensates for inter-individual morphological variations and does not model a physical deformation, the underlying ad-hoc mechanical properties are not related to the actual rheology of the organ under study.

Using the notation from Eq. (2), we now define the Jacobian matrix of the overall registration function considered at iteration n , and taken at material point X , as:

$$J_n(X) := \frac{\partial R_n}{\partial X}(X) = \frac{\partial (r_n \circ \dots \circ r_1)}{\partial X}(X)$$

Eq. (5) shows that this matrix can be updated after each elementary registration step by first multiplying the previous matrix with the Jacobian matrix of the new elementary deformation r_{n+1} taken at the current position of the considered material point, $(r_n \circ \dots \circ r_1)(X)$.

During the registration assembly, the elastic energy stored in the virtual solid is measured at a set of material points $\{X^i\}_i$, evenly distributed among the initial source data, and set to “probe” the space distortions induced by the accumulated elementary deformations. To do so, at each iteration n the Green-Lagrange strain tensor ϵ_n is derived from the above mentioned Jacobian matrix, as $\epsilon_n = (J_n^T J_n - I)/2$, and the stress tensor σ_n is related to the strain tensor ϵ_n by a linear constitutive equation $\sigma_{nij} = D_{ijkl} \epsilon_{nkl}$.

The potential elastic energy generated by the deformation R_n at control point X^i can be computed using the Total Lagrangian formulation, as:

$$W_n^i = \int_{V^i} \epsilon_n : \sigma_n dX$$

where V^i is the volume in the initial source configuration associated to, or “monitored” by the material point X^i . In order to preserve fast registration computation, the above integral is approximated as $W_n^i \approx |V^i| \epsilon_n(X^i) : \sigma_n(X^i)$, where $|V^i|$ is the measure in the initial configuration of the volume V^i associated to the material control point X^i .

The sum of the contributions of all control points $\{X^i\}_i$ gives an approximation of the total potential elastic energy stored in the deformed source space at iteration n , as $W_n = \sum_i W_n^i$. This measure is taken into account at each deformation step to select, among all possible elementary deformations, the registration function r_n which offers the best ratio between registration energy decrease and elastic energy increase.

To this end, at each iteration n and for each candidate deformations r_n the associated registration energy decrease $\Delta E_n > 0$ is computed as $\Delta E_n = E(R_{n-1}) - E(r_n \circ R_{n-1})$. The change in potential elastic energy ΔW_n associated to each r_n is also computed as $\Delta W_n = W_n - W_{n-1}$ and the following selection algorithm is applied.

1. Among all candidate elementary deformations, only the deformations r_n leading to a relative registration gain greater than the stop threshold T_E are considered, i.e. the ones for which $\Delta E_n/E_{n-1} > T_E$.
2. Among those, if deformations such as $\Delta W_n < 0$ can be found, the one yielding the highest registration energy decrease is chosen.
3. Otherwise, the deformation r_n having the highest $\Delta E_n/\Delta W_n$ ratio is applied.

Rule 1 merely implements the stop criterion mentioned above. If no elementary registration function satisfying this condition can be found, the iterations stop. Rule 2 favors space “decompression” if it goes with satisfactory registration enhancement. Finally Rule 3 is applied in the most general case to select the deformation offering the best trade-off between space distortion and registration gain.

It is important to emphasize here that although the elastic grid initial state is restored prior to each elementary registration step (Eulerian approach), the mechanical energy terms W_n computed above keep track of all the deformations accumulated in R_n at step n (Lagrangian formulation). They are therefore appropriate for measuring space distortion as the registration advances.

As a relationship between inter-individual shape variations and mechanical behavior could not be determined, the simple St. Venant–Kirchhoff mechanical framework was selected. An isotropic soft compressible material using an empirically set Poisson's ratio $\nu = 0.2$ and a Young's modulus $E = 1$ Pa was implemented. The value of E has no effect on the registration sequence as it is not the sum but the ratio between registration and elastic energy that conditions each elementary deformation.

Finally, in our implementation, 1000 control points are used. The $\{X^i\}_{i=1,\dots,1000}$ are distributed in the bounding box of the source data, on a regular $10 \times 10 \times 10$ grid and all $|V^i|$ are set to 1/1000 of the bounding box volume.

2.7. Multiple structures registration

Prior to mesh registration, a subset of the Atlas mesh nodes needs to be labeled as anatomical features that can be identified within the patient data.

In the case of the femur models (see Sections 4.1 and 4.2) the surface mesh nodes are labeled as “bone surface” as they lie on the cortical surface of the bone, which can easily be extracted from a Computed Tomography (CT) volume.

As for the face model, two families of nodes are defined, labeled “skin” for the exterior surface nodes registered onto the patients' faces skin, and “bone” for the interior surface nodes that correspond to the skull features and need to be registered onto segmented patients' skulls (see Section 4.3). This application example illustrates the capability of our procedure to capture the shapes of multiple anatomical structures modeled within a unique generic FE mesh.

When multiple anatomical structures need to be recovered, the registration of the FE mesh onto the patient data is driven by the minimization of an energy term E^{lab} which measures the fit between the labeled nodes and their corresponding anatomical structures:

$$E^{\text{lab}}(R) = \sum_{l=1}^L \sum_{s \in S_l} d(R(s), D_l)$$

where $l = 1, \dots, L$ are the predefined labels and S_l and D_l are the corresponding source and destination regions, respectively. When multiple labels are defined, a distance map is computed for each D_l subset of the destination data.

The elastic grid deformation is driven by the E^{lab} energy computed over the sets of labeled source nodes and destination surfaces. The unlabeled source points, on the other hand, are embedded within the elastic grid and passively follow the deformation without contributing to the energy term.

3. Mesh repair

Although the elastic registration algorithm described above strongly limits space distortions, the registered mesh may exhibit irregular or low quality elements that need to be untangled before proceeding to FE analysis. Indeed, the non-folding nature of the registration function is a local property ensuring that space orientation is preserved. While this is locally true at every point in space, the property does not hold when considering finite structures such as element edges and, as a consequence, irregular elements may appear after the application of the smooth and non-folding elastic deformation to the initially regular Atlas mesh.

Element **regularity** is assessed by considering the mapping $F: \xi \rightarrow x$, between the element parent (or reference) coordinates system (ξ_1, ξ_2, ξ_3) and the actual element coordinates (x_1, x_2, x_3) . The Jacobian $J(\xi)$ is the determinant of the matrix $\partial F / \partial \xi$. Its value

represents, at a given reference point ξ , the local volume transformation between the parent and actual element configuration.

An element is said to be regular if $J(\xi) > 0$ for all ξ in the parent configuration. Element regularity is usually assessed by considering the value of J at specific points such as the integration points or the element nodes (Hughes, 1987). We call a node \mathbf{n} “irregular” if it has a negative Jacobian $J_{\mathbf{n}}^e$, computed in element e . A mesh is said “regular” if all the elements, and hence all nodes, are regular. An irregular mesh is not suitable for Finite Element analysis as the singularity of the mapping F leads to modeling inconsistency.

Element **quality**, on the other hand, is a measure of the conformity of its shape, which reflects the evenness of the discretization of the modeled domain. There is a great variety of quality measures and their relevance is dependent on the considered element type and computations to be carried out (Field, 2000; Shewchuk, 2002).

Given the fact that fast and robust tetrahedral discretization solutions are already available, such as the widely used TetGen software, or have been proposed in the literature (Molino et al., 2003; Frey, 2004; Alliez et al., 2005), we focus our work on hexahedral-dominant meshes and demonstrate our mesh generation technique using a popular quality measure, well suited for hexahedrons and wedges: the Jacobian ratio (JR) (Knupp, 2000).

The JR is defined for a given node \mathbf{n} considered within element e . Its value is the ratio between the Jacobian at node \mathbf{n} in e , and the maximal nodal Jacobian value in element e , $J_{\max}^e = \text{Max}_{\mathbf{m} \in e} \{J_{\mathbf{m}}^e\}$, thus:

$$\text{JR}_{\mathbf{n}}^e = \frac{J_{\mathbf{n}}^e}{J_{\max}^e}$$

By comparing, the Jacobian value of each node with the maximal value within the considered element, the JR gives an indication of the contribution of a node to the overall element distortion, as opposed to local nodal distortion between the parent and actual configuration measured by the Jacobian alone. Element quality with respect to parent configuration is proportional to the JR value, ranging from 0 to 1.

Mesh repair is a twofold process: first all the elements in the mesh are inspected and, if necessary, their nodes' positions are adapted so as to recover **regularity**; then a second relaxation procedure is carried out on the mesh if the achieved **quality** levels are unacceptable. Both steps can affect the nodes positions and alter the organ surface representation achieved after the elastic registration step.

In the absence of a formal proof that an acceptable mesh configuration exists and can be found by the relaxation procedure in all situations, the aim of this study is to evaluate the performance of MMRRep on a database of diverse organ shapes and clinical use cases. The results described in Sections 4.1–4.3 suggest that the smoothness of the elastic registration strongly limits spatial distortion making it possible to recover both mesh regularity and quality without interfering with the prior surface registration, by applying small displacements to a limited subset of the surface nodes.

A large number of nodes in a mesh can make the repair procedure computationally prohibitive. This complexity can be greatly reduced using a local relaxation strategy, i.e. grouping all irregular or poor quality nodes into “regions” defined in such a way that nodal corrections applied inside a region leave the outside mesh configuration unchanged. This local repair strategy makes it possible to perform all relaxation procedures independently on each repair region identified within the mesh. It also decreases significantly the computational complexity as the number of degrees of freedom to be considered in a region is usually small.

The untangling of an irregular region R consists in finding a configuration where all the Jacobians in R , $\{J_j\}_{j \in R}$, are positive. This nodal relaxation can be formulated as a maximization procedure driven by a “regularity energy” E_R . E_R is expressed as the sum of

all Jacobians within R affected by a penalty function φ_k , of strength controlled by an index k , giving:

$$E_R = \sum_{j \in R} \varphi_k(J_j)$$

where $\varphi_k(t) = 1 - \exp(-kt)$. As the parameter k is increased during the optimization process, the influence of negative values overbalances the positive ones thus favoring a solution where all Jacobians in the sum are positive.

The aim of the quality maximization, on the other hand, is to find a configuration where all the JRs in a region R are above a pre-defined level JR_{min} . The associated “quality energy” E_Q is defined as the sum of the JRs within R weighted by a penalty function ψ_k , thus:

$$E_Q = \sum_{j \in R, e \in R} \psi_k(JR_j^e)$$

where $\psi_k(t) = 1 - \exp(k(JR_{min} - t))$. Similarly to the regularization process, the penalty parameter k is used to find a solution where all JRs contributing to the sum E_Q are above the quality threshold JR_{min} . The value of JR_{min} was chosen in accordance with the quality standard requested by the commercial FE analysis software ANSYS Workbench (ANSYS Inc., USA), i.e. $JR_{min} = 1/30$ (Kelly, 1998).

The initial value of k for both penalty functions φ_k and ψ_k must be carefully chosen. Indeed, given the formulation of the penalty functions, an excessive penalization level may induce strong numerical instabilities in the optimization process. To avoid this issue, the starting value of k is determined by considering the slope of the penalty function at the most penalized energy term (minimal Jacobian during regularization phase; minimal JR during quality optimization phase) and ensuring that it does not exceed a predefined threshold.

Both regularity and quality optimizations are carried out by gradient ascent. Gradients of both E_R and E_Q energies are computed using the centered differences scheme. Assuming unimodality of the local energy function, the maximum search in the direction of ascent is done using the golden section technique (Press et al., 1992) between the current nodal configuration and the configuration obtained after applying the maximal amplitude correction.

Table 1

Registration (s): elastic registration times, in seconds; CT points: quantity of segmented points in CT volumes; Mean, Max, σ (mm): surface representation mean and maximal error, standard deviation, in millimeters.

Patient	1	2	3	4	5
Registration (s)	25	36	42	26	32
CT points	10,930	25,980	22,924	20,065	17,886
Mean err. (mm)	0.3	0.4	0.3	0.4	0.3
Max err. (mm)	5.4	6.6	5.2	6.0	5.5
σ (mm)	0.5	0.5	0.4	0.5	0.4

Table 2

Regularity, quality (s): mesh repair times for both phases, in seconds; % nodes, nodes/4052: fraction and number of nodes moved by the repair procedure; Mean, Max, σ (mm): mean and maximal nodal displacements, standard deviation, in millimeters.

Patient	1	2	3	4	5
Regularity (s)	2.2	0.6	0.4	0.9	0.5
Quality (s)	3.6	2.6	6.8	0.9	0.7
% nodes	1.0	0.3	0.6	0.5	0.4
Nodes/4052	39	27	23	22	17
Mean disp. (mm)	1.0	0.3	0.2	0.5	0.2
Max disp. (mm)	3.4	2.5	1.1	1.9	0.6
σ (mm)	1.0	0.5	0.2	0.5	0.2

Table 3

Registration (s): elastic registration times, in seconds; points: number of intraoperatively digitized points; Mean, Max, σ (mm): surface representation mean and maximal error, standard deviation, in millimeters.

Patient	1	2	3	4	5
Registration (s)	15	17	20	11	23
Points	1204	1433	1421	1450	1437
Mean err. (mm)	0.3	0.4	0.4	0.3	0.3
Max err. (mm)	2.1	3.4	4.2	2.2	2.7
σ (mm)	0.3	0.4	0.4	0.3	0.3

The amplitude and number of iterations for each mesh region are limited so as to restrict loss of surface representation accuracy. Our implementation allows a maximum of 50 iterations, and nodal displacements at each step have a maximal amplitude of 0.1 mm. The $50 \times 0.1 = 5$ mm maximal displacement can only be achieved by moving a unique node in a constant direction throughout the repair procedure, which seldom occurs. During the experimental validation discussed in Section 4, nodal corrections with mean amplitude less than 1.2 mm and applied to less than 1% of the nodes were sufficient to repair the 60 registered meshes.

The value of maximal nodal correction used here is not universal and must be determined for each field of application according to the dimensions of the modeled domain and maximal tolerance on the representation of its geometry. In our case, the results were consistent with the desired submillimetric mean surface representation accuracy, as shown in Tables 1, 3 and 5.

4. Experimental evaluation

The MMRRep technique was evaluated in three different situations where:

- **complete organ geometry** can be retrieved from the available data, as discussed in Section 4.1;
- only **partial organ geometry** is accessible, as illustrated in Section 4.2;
- **distinct organ features** need to be taken into account by the generated model, as shown in Section 4.3.

The mesh adaptation technique alone is presented here and the discussion about the biomechanical simulations illustrating the three following “use cases” falls out of the scope of this article.

4.1. Complete pre-operative femora CT scans

In this section, we demonstrate our technique in the context of total knee arthroplasty (TKA). The prosthesis placement can be optimized to avoid unsealing or femur fracture using biomechanical modeling and FE analysis of the stresses within the tissues. These tissues are modeled based on the bone mechanical properties, geometry and prescribed loads inferred from the patient morphology and gait (Zalzal et al., 2008).

4.1.1. Mesh registration procedure

The manually assembled right femur Atlas mesh (Couteau et al., 1998) used in this part of the study is composed of 4052 nodes, forming 3018 elements: 2960 hexahedrons and 58 wedges (6 nodes prisms). The elements are organized so as to reflect the bony structure such as the femoral diaphysis cortex which is discretized by a single layer of elements. The principal mesh features are illustrated in Fig. 4.

For the five considered patients a CT scan of the right leg was pre-operatively acquired and a semi-automatic threshold-based

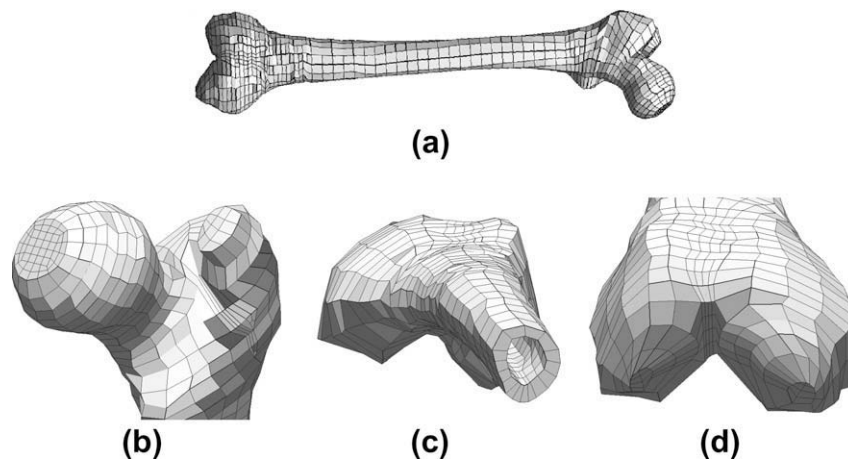


Fig. 4. Femur Atlas mesh from Couteau et al. (1998). (a) Overview. (b) Femur head and great trochanter. (c) Cut out showing the diaphysis cortex layer of elements. (d) Distal condyles.

segmentation procedure was performed by a clinician in order to identify the femur cortical surface. Then the resulting points cloud was rigidly registered onto the Atlas bone surface using an Iterative Closest Point algorithm (Besl and McKay, 1992).

Each elastic mesh registration was carried out as described in Section 2 by taking the source points set S as S_i , the segmented cortical points for patient i , and the destination set D as the Atlas mesh cortex. Once the segmented points were registered onto the surface of the Atlas, the application of the inverse registration function R^{-1} to the generic mesh produced the patient specific FE model. The chosen accuracy level for the inverse computation, as described in Section 2.4, was $\epsilon = 0.1$ mm. Mesh regularity and quality criteria were subsequently analyzed and the mesh repair procedures described in Section 3 were applied to the model.

4.1.2. Results

Table 1 describes the performance of the mesh registration procedure for the five data sets. The surface representation error is computed as the distance between the segmented CT points and the generated patient specific FE mesh surface. The registration times include the direct registration R computation as well as the application of the inverse registration function R^{-1} to the Atlas mesh nodes.

Table 2 gives the performance of the mesh repair procedure carried out on the five deformed meshes. Repair computation times and numbers of corrected nodes are given along with nodal displacements statistics.

The MMRep procedure succeeded in generating a quality femur mesh for all five patients. All computations were carried out in less than 1 min, which is acceptable even in an intraoperative context. The surface representation figures remained unchanged after the application of the mesh repair procedure (see Table 1) as the proportion of displaced nodes did not exceed 1% and most of them were inner nodes which could be freely moved without affecting the mesh surface shape. The reported maximal errors are mainly due to manual segmentation irregularities and lack of local refinement of the Atlas mesh, making it difficult to capture some local shape variations. If necessary, this issue could be solved by further refining the generic mesh.

A sample result of the procedure is shown in Fig. 5 with focus on proximal and distal parts of the automatically generated femur mesh.

Mesh quality distribution measured on the 3018-element full femur Atlas mesh is shown in Fig. 6a and b gives the mean quality

distribution in the five generated patient-specific meshes, along with standard deviations.

The histograms presented in this article, in Figs. 6, 8 and 12, classify the elements into five Jacobian ratio categories. The first interval $[0.03, 0.2]$ lists the elements with a quality level that is acceptable from the point of view of our target application ANSYS, but is usually considered to be “questionable”. Due to the manual assembly process, a small number of questionable elements is present in the Atlas meshes used in this study. As shown by the figures, this proportion only slightly increases after the application of the elastic deformation and mesh repair procedures.

4.2. Partial intraoperative femora digitizations

In this part of the study, we demonstrate the possibility of intraoperative FE mesh generation during a total hip arthroplasty (THA) procedure. As in TKA, FE analysis can be used here to optimize femoral stem placement so as to minimize internal stresses and maximize the implant lifetime (Lengsfeld et al., 2005).

4.2.1. Mesh generation procedure

Due to the surgical procedure limitations, the complete shapes of the patients' femoral heads were not available pre-operatively and each bone geometry was acquired intraoperatively by sliding a calibrated pointer on the cortical surface of the partially exposed hip. The pointer position was tracked in space by means of an optical localization system (Polaris, NDI, Canada) and the position of its tip was recorded continuously.

The initial positions of the recorded points with respect to the Atlas model were computed using the correspondence between anatomical landmarks such as the knee center and the piriformis fossa, localized in patient space using the pointer, and defined in the Atlas space by an expert operator. Fig. 7 shows the similar and very localized distributions of the digitized points for all 5 patients with respect to an approximative surface model of the proximal femur (3 right and 2 left hips were included in this study).

The hip Atlas mesh used here is the upper part of the complete right femur Atlas presented in Section 4.1. After truncation, the right hip model was mirrored with respect to the sagittal plane to produce the left hip generic model. It is composed of 2105 nodes, forming 1738 hexahedrons and 16 wedges organized so as to reflect the hip principal mechanical structures such as the femoral head and neck.

Mesh registration and repair was carried out as described in Section 4.1, taking S as S_i , the digitized points cloud for patient i ,

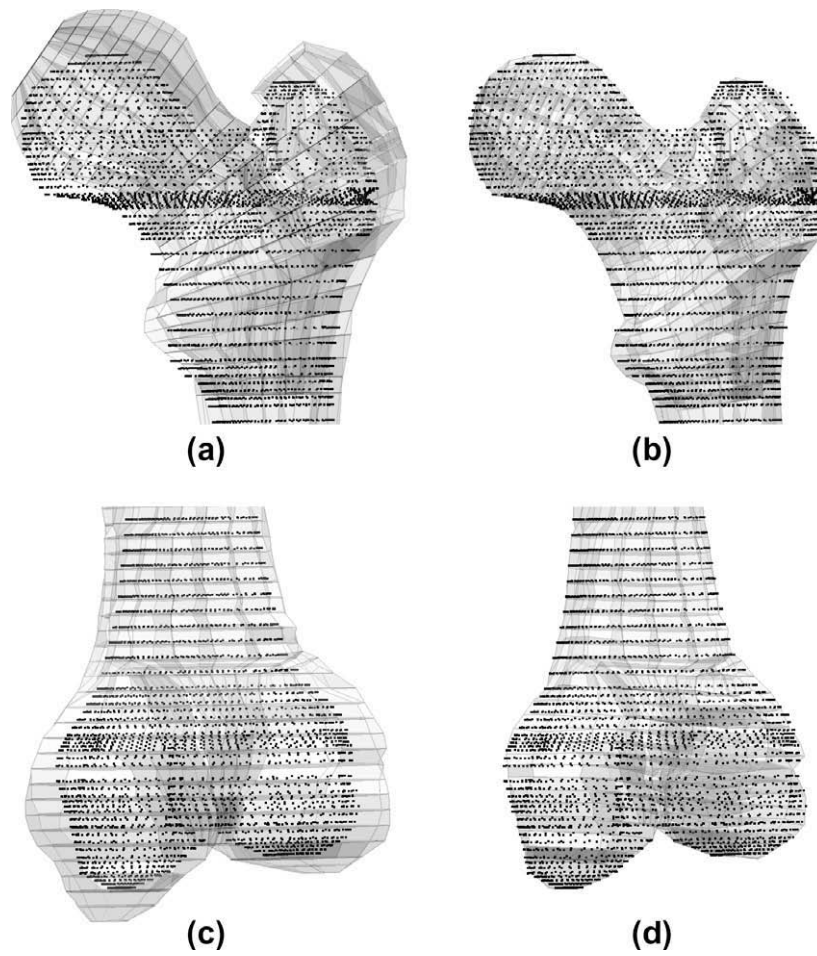


Fig. 5. Sample patient specific mesh. Proximal epiphysis: (a) segmented patient data (black dots) and Atlas mesh; (b) patient specific mesh fitting the femur head surface. Distal epiphysis: (c) segmented patient data (black dots) and Atlas mesh; (d) patient specific mesh fitting the condyles.




Fig. 6. Mesh quality statistics for Atlas (a) and patient-specific (b) full femur models.




Fig. 7. Intraoperatively acquired point clouds for the five considered patients (black dots). To localize the intraoperatively accessible bone region, a surface mesh of the femoral head is showed with each data set (wireframe surface).

and D as the Atlas hip bony surface. Each patient specific mesh was created by applying the inverse R^{-1} of the registration function computed with an accuracy level $\epsilon = 0.1$ mm, as described in Section 2.4.

4.2.2. Results

The two result tables below are similar to those included in Section 4.1. Table 3 gives the performance of the mesh registration procedures carried out on the five data sets, the surface representation error being, this time, computed by considering the distance between all digitized points and the generated patient specific hip surface.

Table 4 presents the performance of the mesh repair procedures for the five meshes. The repair times are given in seconds. In the case of patient 4, the patient mesh produced by the elastic registration was already regular.

Mesh quality distribution measured on the 1754 elements hip Atlas mesh is shown in Fig. 8a and b gives the mean quality distribution in the five generated patient-specific meshes, along with standard deviations.

Table 4

Regularity, quality (s): mesh repair times for both phases, in seconds; % nodes, nodes/2105: fraction and number of nodes moved by the repair procedure; Mean, Max, σ (mm): mean and maximal nodal displacements, standard deviation, in millimeters.

Patient	1	2	3	4	5
Regularity (s)	0.4	0.1	0.5	0	0.4
Quality (s)	0.5	0.4	0.8	0.4	0.4
% nodes	0.3	0.05	0.7	0.05	0.4
Nodes/2105	6	1	14	1	8
Mean disp. (mm)	1.2	0.06	0.6	0.04	0.7
Max disp. (mm)	3.0	0.06	2.8	0.04	2.6
σ (mm)	1.0	0.0	0.9	0.0	0.9

As in the previous illustrative case, the MMRep algorithm successfully worked for all five patients. All patient specific meshes were generated in less than 25 s and had submillimetric surface representation accuracy which remained unchanged by the repair procedures affecting less than 0.7% of the nodes. The maximal surface representation errors reported here are due to lack of refinement in the template mesh but also to the presence of noise in the digitized point clouds. Indeed, if the hand-held digitization pointer is lifted from the bony surface during the hip surface acquisition process, erroneous points can be recorded. These outliers are eventually averaged out during the elastic registration process, but their presence is revealed by the surface representation error measures.

This second example demonstrates the adequacy of the MMRep technique in situations where only a fraction of the organ anatomy is known prior to modeling. The elastic registration process, thanks to an a priori knowledge about the organ of interest carried by the Atlas mesh, makes it possible to generate a FE model with high surface representation accuracy in the digitized zones and an approximate yet realistic organ shape in regions where no data is available. The repair phase ensures that the produced model meets the required quality standard and is suitable for FE analysis.

4.3. Skin and bone face modeling

This last use case demonstrates the application of the MMRep procedure to patient specific FE mesh generation in the context of orthognathic surgery, where FE analysis helps predict the consequences of the intervention on the patient's features and facial expressions by simulating the effects of the repositioning of the jaw, maxillary or malar bones (Chabanas et al., 2002, 2003; Luboz et al., 2005).




Fig. 8. Mesh quality statistics for Atlas (a) and patient-specific (b) hip models.

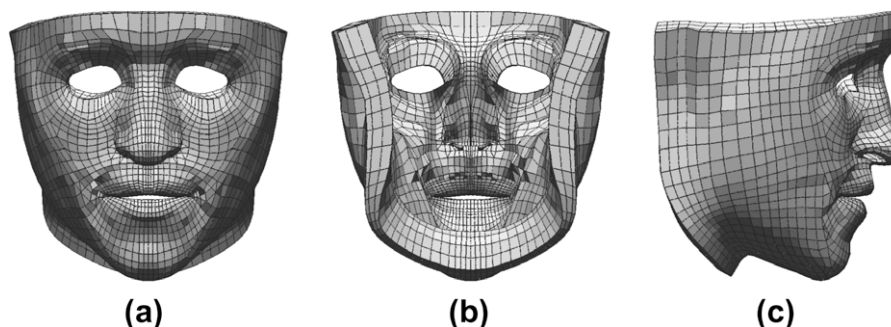


Fig. 9. Face Atlas mesh from Nazari et al. (2008). (a) Anterior view. (b) Posterior view. (c) Lateral view.




Fig. 10. (a) Skin and skull surfaces reconstruction from a sample patient CT volume. (b) Generic face mesh rigidly registered with the patient data.

4.3.1. Mesh registration procedure

In this application a manually assembled 3-layers mesh developed by Nazari et al. (2008), shown in Fig. 9, is used to model the face muscles and fat. It is made of 8746 nodes forming 6030 hexahedrons and 314 wedges. As we wish to model the interaction between face bones, muscles and features, the patient specific FE model must fit both the skin and skull reconstructed from each CT volume. To this end the Atlas outer and inner layer nodes are labeled “skin” and “bone”, respectively, and the multi-labels formulation of the elastic registration discussed in Section 2.7 is used. The inside Atlas mesh nodes, defining the inner layers of the face tissues, are unlabeled and follow the overall elastic deformation driven by the registration of the nodes labeled “skin” and “bone”.

For the 50 patients included in this study (data provided by the MAP5 Laboratory, University of Paris V), the bone and skin layers were segmented in the CT volumes using the Hounsfield scale and the resulting surfaces were reconstructed (Tilotta, 2008) and oriented along the anatomical axes in accordance with the generic face model. The Atlas mesh was aligned on the patient data using the nose tip position. Fig. 10 shows sample patient skin and bone surfaces along with the translated Atlas model before elastic registration.

The generic face mesh represents a subset of the complete patient head data, therefore, as discussed in Section 2.5, the computed registration R is the one that fits the labeled Atlas mesh nodes to their corresponding destination skin or bone surfaces, producing the transformation that has to be applied to the generic mesh to specialize it for the specific patient. From the computational point of view, as a distinct set of distance maps has to be computed for each patient’s skin and skull destination surfaces, the mesh generation process requires more time than in the previously described cases. Yet, in a pre-operative simulation scenario, computational delays are less critical than in an intraoperative FE analysis context.

4.3.2. Results

Unlike the previous use cases, this section does not give the detail of MMRep performance figures for each of the 50 patients. Instead, Table 5 summarizes the overall performance of our technique by presenting the mean registration speed, mesh repair cost and surface reconstruction accuracy. The surface representation errors shown here are the final measures performed after the repair procedure has been applied to the mesh. The nodal displacements amplitudes were evaluated separately for bone and skin layers.

The impact of the mesh repair procedure is very limited as it only affects an average of 2% of the 1715 bone nodes and 0.3% of

Table 5

Elastic mesh registration, regularization and quality optimization times, in seconds. For both bone and skin layers: final surface representation mean errors, number of nodes corrected by the repair procedure in each layer and nodal displacements amplitudes, in millimeters.

	Mean	Max	σ
Elastic registration (s)	32	96	19
Regularity (s)	28	50	12
Quality (s)	4	21	2
Bone			
Surface err. (mm)	0.6	13.8	0.5
Moved nodes/1715	39.8	105	23.9
Displacements (mm)	0.8	3.8	0.5
Skin			
Surface err. (mm)	0.4	23.5	0.4
Moved nodes/2180	7.6	45	10.3
Displacements (mm)	0.3	3.1	0.2

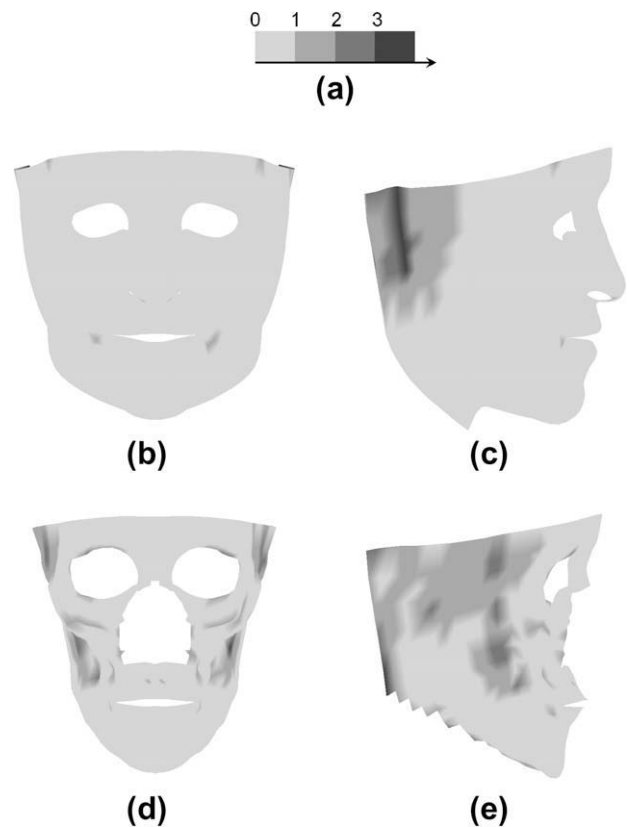


Fig. 11. Face mesh registration mean errors represented as color maps. (a) Error maps color code in millimeters. Skin layer: (b) anterior view; (c) lateral view. Bone layer: (d) anterior view; (e) lateral view.

the 2180 skin nodes in the mesh. As for the surface representation, submillimetric mean accuracy is achieved for both layers.

The reported maximal errors are located in areas where the Atlas mesh lacks refinement. These regions clearly appear in Fig. 11 where the surface representation mean errors computed over the 50 cases are displayed as error maps on the initial Atlas bone and skin layers. The lightest areas represent mean errors between 0 and 1 mm and the darkest areas errors above 3 mm.

Fig. 11 shows that the maximal skin layer errors are located around the ears which are clearly absent from the generic mesh, as can be seen in Fig. 9c. As for the bone layer, the maximal errors appear near the sphenoid bone and the zygomatic process as both




Fig. 12. Mesh quality statistics for Atlas (a) and patient-specific (b) face models.

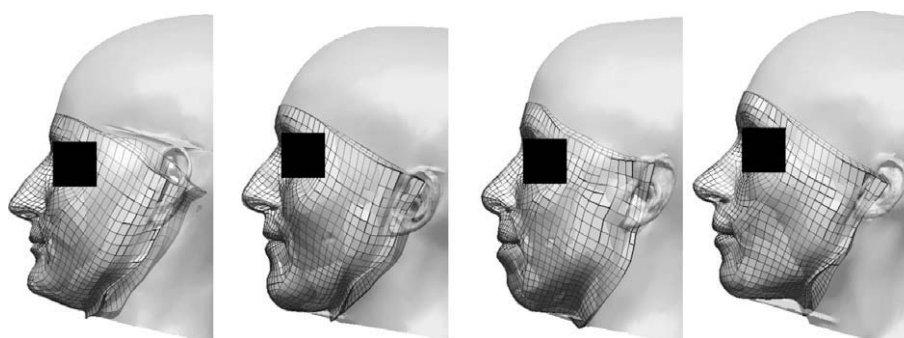


Fig. 13. Four examples of the generated face models. For clarity only the skin surface segmented from the CT volume is shown here although the produced FE models also fit to the underlying skull surface.

regions have a very approximative representation in the Atlas mesh.

Mesh quality distribution measured on the 6344 elements face Atlas mesh is shown in Fig. 12a and b gives the mean quality distribution in the 50 generated patient-specific meshes, along with standard deviations.

In all cases, the MMRep algorithm was able to generate a patient specific bi-layer mesh suitable for FE analysis within a couple of minutes. The produced meshes exhibited submillimetric mean surface representation accuracy on both skin and bone layers, with larger errors localized around features absent or ill-defined in the generic mesh.

Fig. 13 shows four examples of meshes fitted onto distinct patient morphologies. The Atlas face mesh used here was constructed based on a unique prognathic³ patient, yet it was successfully used to model both prognathic and retrognathic⁴ patients, which demonstrates the versatility of the registration procedure.

Finally, Fig. 14 shows 25 thumbnails of registered face meshes demonstrating the variety of clinical cases embraced by the study. The two upper rows show retrognathic cases, the middle row average patients and the two lower rows prognathic morphologies.

5. Discussion and conclusions

A fast and automatic mesh generation procedure – the MMRep algorithm – has been presented, based upon the elastic registration of a generic, or “Atlas”, mesh towards patient specific structures, coupled with a robust mesh repair technique which ensures that

element quality standards are met and FE analysis can safely be carried out on the resulting domain discretization.

To our knowledge this is the first evaluation of a FE mesh registration technique carried out on a wide range of 60 clinical data sets, illustrating three distinct use cases, raising both pre- and intraoperative biomechanical modeling issues and relying on fully or partially available patient data. In all situations the MMRep technique automatically generated a patient specific quality mesh within minutes, which strongly contrasts with time-consuming manual mesh assembly procedures involving a human expert operator.

In all studied cases the regularity of the elastic deformation preserved the mesh elements from excessive distortions and the repair algorithm was able to find a suitable nodal configuration while maintaining a satisfactory surface representation accuracy. Only a small fraction, less than 1%, of the mesh nodes positions needed to be corrected by submillimetric displacements.

Furthermore, the elastic registration formulation made it possible to compute a mesh deformation driven by multiple anatomical structures or sub-structures, as shown in Section 4.3. This feature could easily be used in the femur model generation scenarii, Sections 4.1 and 4.2, for example for distinction between cortical and spongy bone layers, which have distinct mechanical properties.

This study shows that the proposed elastic registration technique is well-suited to the addressed meshing problem: computational times are short and submillimetric surface representation accuracy is achieved. However the MMRep procedure could be supported by any elastic registration algorithm (Beg et al., 2005; Vercauteren et al., 2007; Rueckert et al., 2006), provided that the following properties are ensured:

³ Having the jaws projecting beyond the upper part of the face.

⁴ Having a mandible located posterior to its normal position.




Fig. 14. Sample of 25 registered face meshes. For each patient, the transparent skin surface mesh is shown along with the registered FE mesh. For clarity the skulls were omitted.

- A smooth deformation field, ideally C^1 , can be estimated by the registration algorithm within reasonable computational times.
- Non-folding and bijection of the registration are ensured.
- Accurate registration inverse can be computed if the patient data is registered onto the Atlas and the inverse deformation is applied to the generic mesh in order to make it specific, as described in Sections 4.1 and 4.2.
- The simultaneous registration of different structures, if required by the biomechanical modeling, can be computed as shown in Section 4.3.

The Atlas based approach presented here relies on the definition of a generic model of the target organ with the desired elements layout and possibly some identified sub-structures of interest. This modeling effort only needs to be done once but the resulting Atlas mesh must be carefully designed so that the anatomical features of interest can be properly registered to their patient specific counterparts.

The tests carried out on our database suggest algorithm robustness although no formal proof of convergence could be given. In

extreme cases where strongly distorted pathological organs diverge from the average Atlas shape, mesh regularity can be lost and the repair strategies proposed here may fail. This issue can be overcome by working with a pool of Atlas meshes that represent the main deformation classes likely to be encountered. Such Atlas variations can be easily generated by successfully applying the MMR procedure to a representative case and using the resulting quality mesh as a starting point for subsequent mesh registrations of similar configurations.

A number of enhancements to the presented MMR technique can be foreseen. Breaking the sequential registration-repair scheme, the algorithm could benefit from the integration of the repair process within the elastic registration computation itself. This could be implemented either as a new energy term replacing the ad-hoc mechanical formulation controlling the deformation regularity, or as a per-iteration post-processing callback which, although more straightforward, would unfortunately raise the issue of the convergence of the elastic registration algorithm.

Future works also include the broadening of our mesh repair approach by taking into account other element types such as

pyramids and tetrahedrons, as well as other quality criteria, such as the face warping factor (Robinson and Haggemacher, 1982) that measures each element's face nodes coplanarity. Stronger constraints on the evenness of the generated meshes should enhance the numerical stability of the FE analysis carried out.

Finally, we can imagine an ideal FE mesh generation algorithm that performs mesh registration and repair directly in patient 3D image space. This fully integrated organ modeling tool could be achieved by replacing the distance based elastic registration formulation used here by an appropriate segmentation energy suitable for the considered imaging modality.

Acknowledgements

This work was carried out within the French national research project ANR/TecSan HipSurf. The authors wish to thank Mr. Yves Rozenholc, from the Laboratoire de Mathématiques Appliquées (MAP5, UMR CNRS, University Paris 5) for the head CT scan volumes, Mr. Pascal Swider from the Toulouse biomechanics laboratory for the complete femora CT scans and Mr. Lionel Carrat (Praxim, France) for providing us with the intraoperative hip digitization data sets.

References

- Alliez, P., Cohen-Steiner, D., Yvinec, M., Desbrun, M., 2005. Variational tetrahedral meshing. *ACM Transactions on Graphics* 24, 617–625.
- Beg, M.F., Miller, M.I., Trounev, A., Younes, L., 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision* 61, 139–157.
- Belytschko, T., Liu, W.K., Moran, B., 2006. *Nonlinear Finite Elements for Continua and Structures*. Wiley.
- Benzley, S.E., Perry, E., Merkle, K., Clark, B., Sjaardema, G.D., 1995. A comparison of all hexagonal and all tetrahedral finite element meshes for elastic and elastoplastic analysis. In: *Proceedings of the 4th International Meshing Roundtable*, pp. 179–191.
- Besl, P.J., McKay, N.D., 1992. A method for registration of 3d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 239–254.
- Castellano-Smith, A., Hartkens, T., Schnabel, J., Hose, D., Liu, H., Hall, W., Truwit, C., Hawkes, D., Hill, D., 2001. Constructing patient specific models for correcting intraoperative brain deformation. *Lecture Notes in Computer Science* 2208, 1091–1098.
- Chabanas, M., Payan, Y., 2000. A 3d finite element model of the face for simulation in plastic and maxillo-facial surgery. *Lecture Notes in Computer Science* 1935, 1068–1075.
- Chabanas, M., Marécaux, C., Payan, Y., Boutault, F., 2002. Models for planning and simulation in computer assisted orthognatic surgery. *Lecture Notes in Computer Science* 2489, 315–322.
- Chabanas, M., Luboz, V., Payan, Y., 2003. Patient specific finite element model of the face soft tissue for computer-assisted maxillofacial surgery. *Medical Image Analysis* 7, 131–151.
- Couteau, B., Hobatho, M., Darmana, R., Brignola, J., Arlaud, J., 1998. Finite element modelling of the vibrational behaviour of the human femur using CT-based individualized geometrical and material properties. *Journal of Biomechanics* 31, 383–386.
- Couteau, B., Payan, Y., Lavallée, S., 2000. The mesh-matching algorithm: an automatic 3d mesh generator for finite element structures. *Journal of Biomechanics* 33, 1005–1009.
- Field, D.A., 2000. Qualitative measures for initial meshes. *International Journal for Numerical Methods in Engineering* 47, 887–906.
- Frey, P., 2004. Generation and adaptation of computational surface meshes from discrete anatomical data. *International Journal for Numerical Methods in Engineering* 60, 1049–1074.
- George, P.L., Borouchaki, H., Laug, P.L., 2002. An efficient algorithm for 3d adaptive meshing. *Advances in Engineering Software* 33, 377–387.
- Gibson, A., Riley, J., Schweiger, M., Hebden, J., Arridge, S., Delpy, D., 2003. A method for generating patient-specific finite element meshes for head modelling. *Physics in Medicine and Biology* 48, 481–495.
- Grosland, N.M., Bafna, R., Magnotta, V.A., 2009. Automated hexahedral meshing of anatomic structures using deformable registration. *Computer Methods in Biomechanics and Biomedical Engineering* 12, 35–43.
- Hughes, T., 1987. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Dover Publications.
- Huiskes, H., 1980. Stress analyses of implanted orthopaedic joint prostheses for optimal design and fixation. *Acta Orthopaedica Belgica* 46, 711–727.
- Jaspers, P., De Lange, A., Huiskes, H., Van Rens, T., 1980. The mechanical function of the meniscus experiments on cadaveric pig knee-joint. *Acta Orthopaedica Belgica* 46, 663–668.
- Kelly, S., 1998. Element shape testing. In: *Ansys Theory Reference*, ninth ed. (Chapter 13).
- Knupp, P., 2000. Achieving finite element mesh quality via optimization of the jacobian matrix norm and associated quantities. Part II – a framework for volume mesh optimization and the condition number of the jacobian matrix. *International Journal for Numerical Methods in Engineering* 48, 1165–1185.
- Kwok, W., Chen, Z., 2000. A simple and effective mesh quality metric for hexahedral and wedge elements. In: *Proceedings of the 9th International Meshing Roundtable*, pp. 325–333.
- Lengsfeld, M., Burchard, R., Gunther, D., Pressel, T., Schmitt, J., Leppek, R., Griss, P., 2005. Femoral strain changes after total hip arthroplasty – patient-specific finite element analyses 12 years after operation. *Medical Engineering and Physics* 27, 649–654.
- Liao, S., Tong, R., Wang, M., Dong, J., 2005. Rapidly generate lumbar spine volume mesh. In: *Ninth International Conference on Computer Aided Design and Computer Graphics*. IEEE Computer Society, pp. 345–350.
- Luboz, V., Pedrono, A., Ambard, D., Boutault, F., Payan, Y., Swider, P., 2004. Prediction of tissue decompression in orbital surgery. *Clinical Biomechanics* 19, 202–208.
- Luboz, V., Chabanas, M., Swider, P., Payan, Y., 2005. Orbital and maxillofacial computer aided surgery: patient-specific finite element models to predict surgical outcomes. *Computer Methods in Biomechanics and Biomedical Engineering* 8, 259–265.
- Molino, N., Bridson, R., Teran, J., Fedkiw, R., 2003. A crystalline, red green strategy for meshing highly deformable objects with tetrahedra. In: *Proceedings of the 12th International Meshing Roundtable*, Sandia National Laboratories. Springer-Verlag, pp. 103–114.
- Nazari, M., Payan, Y., Perrier, P., Chabanas, M., Lobos, C., 2008. A continuous biomechanical model of the face: a study of muscle coordinations for speech lip gestures. In: *Proceedings of the 8th International Seminar on Speech Production – ISSP'08*, pp. 321–324.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 1992. *Numerical Recipes in C*. Cambridge University Press.
- Robinson, J., Haggemacher, G.W., 1982. *Element warning diagnostics*. Finite Element News.
- Rueckert, D., Aljabar, P., Heckemann, R.A., Hajnal, J.V., Hammers, A., 2006. Diffeomorphic registration using b-splines. *Medical Image Computing and Computer-Assisted Intervention MICCAI*, 702–709.
- Saito, T., Toriwaki, J., 1994. New algorithms for euclidean distance transformations of an n-dimensional digitized picture with applications. *Pattern Recognition* 27, 1551–1565.
- Shewchuk, J., 2002. What is a good linear element? Interpolation, conditioning, and quality measures. In: *Proceedings of the 11 International Meshing Roundtable*, Sandia National Laboratories. Springer-Verlag, pp. 115–126.
- Shim, V., Pitto, R., Streicher, R., Hunter, P., Anderson, I., 2007. The use of sparse CT datasets for auto-generating accurate fe models of the femur and pelvis. *Journal of Biomechanics* 40, 26–35.
- Si, H., 2006. On refinement of constrained delaunay tetrahedralizations. In: *Proceedings of the 15th International Meshing Roundtable*, pp. 509–528.
- Si, H., Gaertner, K., 2005. Meshing piecewise linear complexes by constrained delaunay tetrahedralizations. In: *Proceedings of the 14th International Meshing Roundtable*, pp. 147–163.
- Sigal, I.A., Hardisty, M.R., Whyne, C.M., 2008. Mesh-morphing algorithms for specimen-specific finite element modelling. *Journal of Biomechanics* 41, 1381–1389.
- Taddei, F., Pancanti, A., Viceconti, M., 2004. An improved method for the automatic mapping of computed tomography numbers onto finite element models. *Medical Engineering and Physics* 26, 61–69.
- Tilotta, F., 2008. *Contribution à la reconstitution faciale en médecine légale: proposition d'une nouvelle méthode statistique*. Ph.D. thesis, Université Paris Sud 11, France (2/7/08).
- Vercateren, T., Pennec, X., Perchant, A., Ayache, N., 2007. Non-parametric diffeomorphic image registration with the demons algorithm. *Proceedings of Medical Image Computing and Computer-Assisted Intervention MICCAI* 4792, 319–326.
- Viceconti, M., Davinelli, M., Taddei, F., Cappello, A., 2004. Automatic generation of accurate subject-specific bone finite element models to be used in clinical studies. *Journal of Biomechanics* 37, 1597–1605.
- Wittek, A., Miller, K., Kikinis, R., Warfield, S.K., 2007. Patient-specific model of brain deformation: application to medical image registration. *Journal of Biomechanics* 40, 919–929.
- Zalzal, P., Papini, M., Backstein, D., Gross, A.E., 2008. Anterior femoral notching during total knee arthroplasty: a finite element analysis. *Journal of Bone and Joint Surgery, British* 90-B.