



HAL
open science

Geometry description and mesh construction from medical imaging

Michele Giuliano Carlino, Philippe Ricka, Minh Son Phan, Silvia Bertoluzza,
Micol Pennacchio, Giuseppe Patanè, Michela Spagnuolo

► **To cite this version:**

Michele Giuliano Carlino, Philippe Ricka, Minh Son Phan, Silvia Bertoluzza, Micol Pennacchio, et al.. Geometry description and mesh construction from medical imaging. ESAIM: Proceedings and Surveys, 2020, 67, pp.161-177. 10.1051/proc/202067010 . hal-02072342v3

HAL Id: hal-02072342

<https://inria.hal.science/hal-02072342v3>

Submitted on 12 Nov 2019

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

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

GEOMETRY DESCRIPTION AND MESH CONSTRUCTION FROM MEDICAL IMAGING *

MICHELE GIULIANO CARLINO^{1,2}, PHILIPPE RICKA³, MINH PHAN⁴, SILVIA BERTOLUZZA⁵,
MICOL PENNACCHIO⁵, GIUSEPPE PATANÈ⁵ AND MICHELA SPAGNUOLO⁵

Abstract. We present a new method for defining and meshing patient-specific domains from medical images. Our approach is based on an atlas image segmentation technique, and relies on the modular registration algorithm of S. Bertoluzza *et al.* [25]. The mesh of the patient-specific domain is generated by deforming the corresponding mesh on an a priori segmented and meshed reference image (the atlas). Our method aims at automating the process at the interface of medical imaging and numerical simulation, thus reducing the computational cost in those situations where simulations have to be managed on numerous medical images of similar patients.

Résumé. Nous présentons une nouvelle méthode de reconnaissance et de maillage d'un domaine d'intérêt d'une image médicale. Notre approche se base sur une méthode de segmentation à partir d'un atlas, et dépend de la boîte à outils modulaire pour la co-registation d'images développée par S. Bertoluzza *et al.* [25]. Le maillage du domaine spécifique au patient est généré par déformation du maillage correspondant sur une image de référence segmentée a priori (l'atlas). Notre méthode vise à automatiser le processus à l'interface entre l'imagerie médicale et la simulation numérique, avec pour but de réduire le coût de calcul dans les situations dans lesquelles des simulations doivent être faites sur de nombreuses images similaires.

INTRODUCTION

The mathematical modeling and numerical simulation of biophysical phenomena relying on patient-specific data is nowadays recognized as an important tool in the diagnosis of diseases and design of personalized treatment, and the demand for increasingly efficient algorithms for carrying out such tasks is attracting the interest of an ever increasing number of researchers. In this framework, whatever the application considered, one of the first key steps is to retrieve, often from a medical image, the geometrical description of the domain corresponding to the organ of interest on which the simulation will be run (called *patient-specific domain*), and to build the corresponding mesh. When dealing with the simulation of many instances of a biophysical phenomenon (be it on many different patients, or for one patient at many different times), it is of paramount importance to have

* This paper has been realized in the framework of ERC Project CHANGE, which has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No 694515)

¹ INRIA, Bordeaux Sud-Ouest, France

² IMB, Université de Bordeaux, France

³ IRMA, Université de Strasbourg, France

⁴ Laboratoire IMATH (EA 2134), Université de Toulon, France

⁵ CNR, IMATI "E. Magenes", Italy

efficient algorithms for defining and meshing the domain. Motivated by this practical requirement, our study aims at defining a method that simultaneously provides the geometry description and a mesh of good quality for the patient-specific domain, to be used for the simulation of a biophysical phenomenon (which reduces to solving a certain partial differential equation).

Our approach exploits the features of *atlas-based segmentation*, a well-known technique [30] where the image is segmented by performing a registration with a pre-segmented, labeled image (the “atlas”) that preserves the geometry and topology of the domains of interest, and represents the anatomical structure and their neighborhood relationships. Indeed, this segmentation yields, as a by-product, a map between the two images, which can be used to transfer additional data or information (in our case, a mesh). Given a reference (atlas) image, and a patient-specific image, the procedure consists in the following main steps (see Figure 1):

- Step 1.** Perform a segmentation of the reference (atlas) image, thus obtaining a model description of the geometry, and generate a mesh of the domain of interest. This reference mesh is called *atlas mesh*.
- Step 2.** Perform a co-registration of the patient-specific image with the reference atlas image, thus obtaining a transformation between the two images; in this way we transfer the segmentation to the new image (*atlas-based segmentation*)
- Step 3.** Using the transformation resulting from Step 2, the atlas mesh is deformed to fit the patient-specific domain, thus obtaining a new, patient-specific, mesh.

In the framework that we are interested in, that is, when we need to treat several patient images sharing the same reference atlas image, Step 1 is performed only once and for all, so that its computational cost is not a key issue. Moreover, Steps 1 and 2 are mutually independent, so that one can begin to run the different instances of Step 2 on the different patient images, while still processing Step 1.

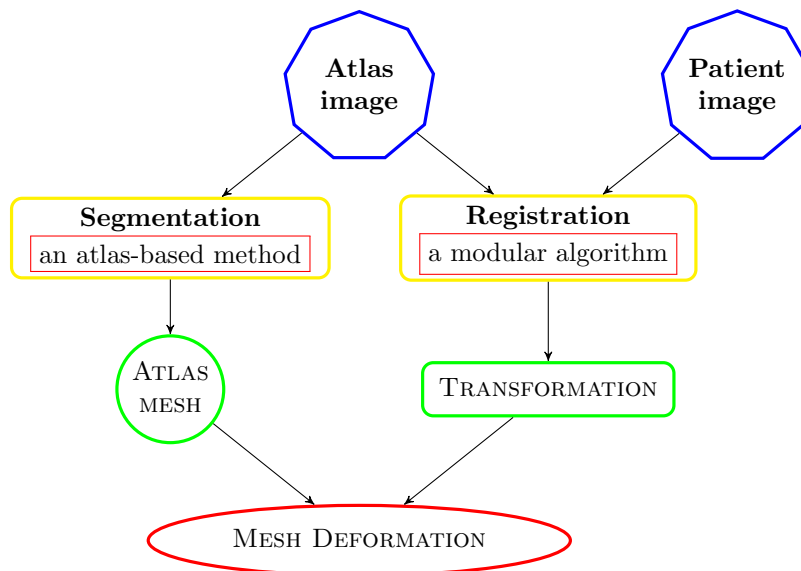


FIGURE 1. Sketch of mesh construction scheme.

The segmentation of the reference image in Step 1 can be performed by any available segmentation technique (in our examples, it is performed “by hand”).

The image registration procedure in Step 2 is carried out by resorting to the modular registration algorithm of S. Bertoluzza *et al.* [25], which splits the registration task in three mutually independent modules by tackling the three different ingredients involved in the registration procedure, namely

- i) the computation of an image metric,
- ii) the rules to carry out image interpolation, extrapolation and differentiation, and
- iii) the definition of the chosen class of parametric image transformations.

The modular structure of the algorithm makes its usage flexible, allowing to independently develop and improve each of the three ingredients.

The paper is structured as follows: Sections 1 and 2 recall the atlas-based image segmentation method and the modular image registration algorithm, respectively. In Section 1, we detail the construction of atlas geometry description and meshes. In Section 2.4, we describe mesh deformations and the resulting creation of patient meshes. Section 3 is devoted to mesh quality analysis and to some simulations solving three model partial differential equations in the patient-related meshes. Section 4 concludes the paper and briefly discusses some open issues and further developments.

1. ATLAS-BASED SEGMENTATION

A monochromatic digital image can be modeled as a matrix (respectively a tensor, for three dimensional images) with positive entries, whose values represent the brightness at the corresponding pixel (respectively voxel). Segmenting such an image consists in assigning a semantic label to each of its pixels (or voxels). There are different approaches to image segmentation, each requiring specific types of image information and constraints on the segmented geometry, and resulting in different methods and algorithms. Just to make some examples, let us mention purely intensity-based classification methods [23, 43, 46], level-set methods [27, 31, 40], snakes or active contours [47], atlas-based segmentation methods [2, 10, 16, 30]. The choice of a suitable method among all those available, relies on different considerations and depends, among other things, on the type of image data that one deals with, on the objects represented by the image, and on the desired output.

We are interested in medical images, such as MR images, CT scans or X-rays, with the main objective to retrieve the geometrical description of an organ of interest (the “patient-specific domain”), and to build a mesh of such a domain in view of some kind of biophysical numerical simulations. Rather than dealing with a specific application, we aim at defining an application independent strategy for the construction of the patient-specific mesh needed for whatever application is at hand. Among different segmentation techniques, we focus our attention on the *atlas-based segmentation*, which takes advantage of the available a priori information on the spatial relation between the different components of the image (which, in our framework, correspond to different anatomical organs or portion of organs), and on their morphometric characteristics. By exploiting this information [21], which can include knowledge about the shape, orientation, continuity, and smoothness of the object to be segmented, atlas-based segmentation can be successful even in the absence of sharp boundaries, and can handle images with no well defined relations between regions and pixels intensities, a situation that can happen in the presence of noise, or when objects of the same texture need to be segmented.

Atlas-based segmentation, has, among other segmentation techniques, a further advantage: together with the segmented image, it provides a map that can be used to transfer any additional information from the reference image (the atlas) to the patient-specific image. In particular, it can be used to transfer to the patient-specific image, information such as a parametrization of the boundary of the domain of interest or the coordinates of the nodes of a finite element mesh for the numerical simulation.

1.1. Atlas and atlas mesh

While our algorithm deals with digital (discrete) images, it is convenient to look at digital images as discrete approximations of continuous images. This allows us to compare images at different resolutions, and gives a natural meaning to mathematical objects such as the gradient of an image. In order to mathematically model the atlas-based segmentation method, let us introduce the following definition, where we normalize images to take values in $[0, 1]$ and where 0 corresponds to black and 1 to white.

Definition 1.1.

- (a) A monochromatic (grey) image is a real valued function $R : \Omega \rightarrow [0, 1]$ where Ω is a subset of \mathbb{R}^d , $d = 2, 3$. Ω is called the domain of image R .
- (b) Given a finite label set $L = \{e_1, e_2, \dots, e_K\}$, a labeling map $A : \Omega \rightarrow L$ is said to be an *atlas*.
- (c) A monochromatic *atlas image* is a couple (R, A) , where R is a monochromatic image and A is an atlas.

Observe that the superposition of A and R yields a segmentation of R into structures identified by the labels of the label set L . Then, an atlas is a segmented image, where different regions, corresponding to the different labels, can be identified. For the atlas to be meaningful, the labeled regions must correspond to the different structures represented in the image.

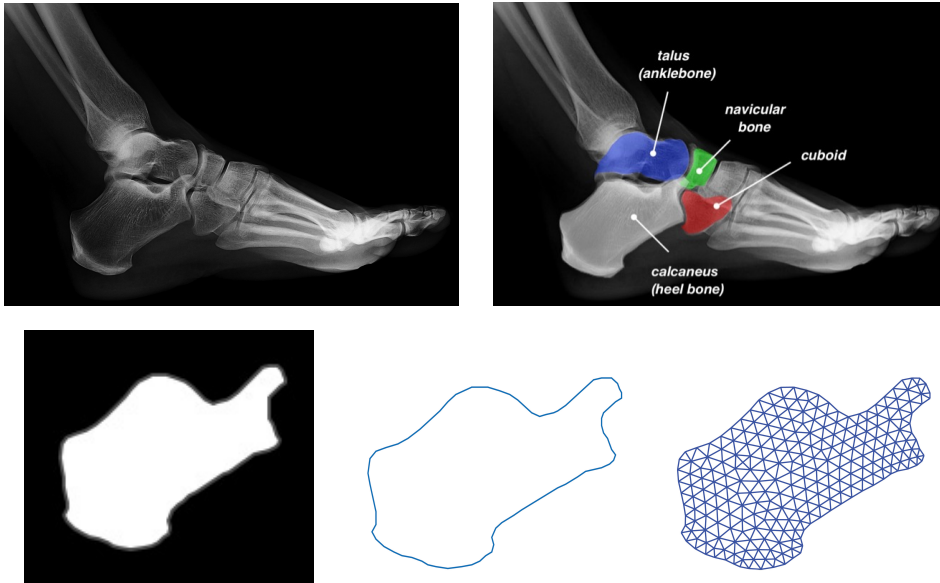


FIGURE 2. On top, on the left the reference (atlas) image of a foot with, on the right a classification of some of the bones and a description of the geometry shape of each bone considered. On the bottom, zooming on the heel bone, from left to right, the simplified black and white atlas image, contour of the region of interest, and the relative reference mesh.

In the domain Ω of the image, we can identify specific subdomains as counterimages of the different labels in L . For each of these subdomains, we can then define an *atlas mesh*, defined as a conforming decomposition of the subdomain into the non overlapping union of triangles or quadrangles (in 3D, of tetrahedra or hexahedra).

The atlas image and mesh can be obtained from a given reference medical image through the following steps:

- (1) segment the medical image thus defining the contours of the regions of interest;
- (2) assign the corresponding labels to pixels belonging to the regions of interest, thus creating the atlas;
- (3) retrieve a set of sample points on the contours of the subdomains identified by the different labels;
- (4) generate the mesh from sample points on the boundaries.

Steps (1) and (2) can be performed by any available mean. We recall that the definition of the atlas is itself an interesting problem, and we refer to [30] for more details. In this paper, the segmentation and labeling of the image are performed by hand using an image editing tool, which results in simplified black and white images (one for each label) where white (resp. black) pixels lie inside (resp. outside) of the anatomical structure of interest. The contours of the regions of interest are then retrieved by applying a level set method to the simplified images. Step (3) is more delicate since a sufficient number of sample points on the contour needs to be retrieved in order to ensure a sufficiently good approximation of the (a priori curved) geometry of the region

of interest. Moreover, the density of sample points directly affects the quality of the mesh, for instance, in the corners of specific domains. There are several available techniques for smoothing the contours (with convolution) and quantifying sample points with respect to approximative curvature. Step (4) is the heaviest step, however, many mesh generators can be found in the literature and, here, we rely on the Delaunay algorithm [12].

In Figure 2, we describe some of the steps of meshing a medical image of a right foot. We are interested in four of the foot bones, namely calcaneus, cuboid, navicular and talus. On top, we see the image on the left and the labeled image (where the four bones have been identified). Observe that the use of atlas-based segmentation allows us to deal with situations where multiple labels are assigned to the same pixel and this represents a further advantage of this segmentation technique (in the present example we observe a superposition between talus, calcaneus and navicular). On the bottom, of the same figure we see the simplified black and white image corresponding to the heel bone, the corresponding contour and triangular mesh.

2. FROM THE ATLAS MESH TO THE PATIENT-SPECIFIC MESH

As mentioned in the introduction, the atlas-based approach reduces the image segmentation problem to an image registration problem. Denoting by \mathbf{Im} the space of monochromatic (continuous) images (Definition 1.1), which we identify with a subset of $L^2(\Omega)$, let $R : \Omega \rightarrow [0, 1]$, $T : \Omega \rightarrow [0, 1]$ be two images in \mathbf{Im} (we assume for simplicity that the two images are defined over the same domain Ω). Registering (or coregistering) the two images means finding, in a given class \mathcal{T} of admissible transformations, a mapping $\theta : \Omega \rightarrow \Omega$ such that $T \circ \theta$ and R are as “close” as possible, as measured by a given “distance” functional $d : L^2(\Omega) \times L^2(\Omega) \rightarrow \mathbb{R}$. In other words, the registration of R and T reduces to the following optimization problem:

$$\text{minimize, over } \mathcal{T}, \text{ the functional } \theta \mapsto d(R, T \circ \theta).$$

The choice of the “distance” functional d and of the class of admissible transformations \mathcal{T} depends on the user’s purpose. For both components, many choices have been studied in the literature, going from the simplest (least squares distance and rigid transformations) to more sophisticated ones [39]. Different choices will lead to different minima, and, consequently to different solutions to the image registration problem. Since in practice we deal with discrete digital images rather than continuous ones, the registration algorithm is also affected by how the discrete data of the pixel values is interpreted, which will lead to different recipes for evaluating quantities such as the point value of an image, or its differential.

In this paper, we restrict ourselves to classes of (non rigid) parametric transformations. For $\bar{\Omega} \in \mathbb{R}^d$, $d = 2, 3$, let us consider a finite dimensional linear subspace of $C^2(\bar{\Omega})$ spanned by a basis $\{\varphi_1(x), \varphi_2(x), \dots, \varphi_M(x)\}$. We denote by \mathcal{T} the class of parametric transformations consisting of all the mappings $\theta_\alpha : \bar{\Omega} \rightarrow \mathbb{R}^d$ of the form

$$\theta_\alpha(x) := x + \Theta(x; \alpha) \tag{1}$$

where, for a parameter vector $\alpha = (\alpha_k^i)_{k,i} \in (\mathbb{R}^M)^d \simeq \mathbb{R}^{Md}$, $\Theta(x; \alpha) \in (C^2(\bar{\Omega}))^d$ is given by

$$\Theta(x; \alpha) = (\Theta_i(x; \alpha))_i, \quad \text{with } \Theta_i(x; \alpha) := \sum_{k=1}^M \alpha_k^i \varphi_k(x). \tag{2}$$

Since the range of $\Theta(\cdot; \alpha)$ is generally not a subset of Ω , to give a meaning to $T \circ \theta_\alpha$, we need to evaluate T also outside Ω . Then we need to introduce an *extension operator* $\mathcal{E} : \mathbf{Im} \rightarrow L^2(\mathbb{R}^d)$ and the optimization problem can then be rewritten as

$$\alpha_{opt} = \arg \min_{\alpha \in \mathbb{R}^{dM}} c(\alpha), \quad \text{with } c(\alpha) := d(R, \mathcal{E}T \circ \theta_\alpha). \tag{3}$$

Problem (3) is a nonlinear, unconstrained problem in \mathbb{R}^{dM} and, to attain a global minimizer instead of local ones, we have to choose an appropriate method. We observe that even in the case that the selected image

distance functional d is convex, the cost function c is not *a priori* a convex function of the unknown α . While locally convergent algorithms like Newton's method and its variants (e.g. Gauss-Newton method and Inexact Gauss Newton method) are usually simple to be implemented and computationally competitive, whenever the initial point α_0 is sufficiently close to some minimizer [22, 29], in our context globally convergent methods, such as steepest descent methods, line search methods, trust-region methods, conjugate-gradient methods, quasi-Newton methods (see [29]), have to be preferred. In our tests we rely on the trust region method, as implemented in the Octave function `fminunc`.

Solving (3) by any gradient-based optimization method requires the (explicit) computation of the gradient of the cost function. Letting $\delta : L^2(\Omega) \rightarrow \mathbb{R}$ be defined as $\delta(T) = d(T, R)$, we recall that, by definition, the Fréchet derivative of δ evaluated at S ($S \in L^2(\Omega)$), is the element $D\delta(S; \cdot)$ of $L^2(\Omega)$ such that for all $\Delta \in L^2(\Omega)$ we have

$$\delta(S + \Delta) = \delta(S) + \int_{\Omega} D\delta(S; x)\Delta(x) dx + o(\|\Delta\|_{L^2(\Omega)}).$$

If we apply the chain rule to the cost function $c(\alpha) = \delta(\mathcal{E}T \circ \theta_{\alpha})$ we obtain the following expression for the gradient $\nabla c(\alpha) = (\partial_{\alpha_1} c(\alpha), \dots, \partial_{\alpha_{dM}} c(\alpha))$:

$$\nabla c(\alpha) = \int_{\Omega} D\delta(\mathcal{E}T \circ \theta_{\alpha}; x) \nabla \mathcal{E}T(\theta_{\alpha}(x)) J_{\alpha} \Theta(x; \alpha) dx, \quad (4)$$

where $\nabla \mathcal{E}T(x) = (\partial_{x_1} \mathcal{E}T(x), \partial_{x_2} \mathcal{E}T(x))$ is the spatial gradient of the (continuous) extended image $\mathcal{E}T$ and $J_{\alpha} \Theta(x; \alpha)$ is the Jacobian of the mapping function Θ with respect to the variable α , that is, for $d = 2$,

$$J_{\alpha} \Theta(x, \alpha) = \begin{pmatrix} \partial_{\alpha_1} \Theta_1(x, \alpha) & \cdots & \partial_{\alpha_{dM}} \Theta_1(x, \alpha) \\ \partial_{\alpha_1} \Theta_2(x, \alpha) & \cdots & \partial_{\alpha_{dM}} \Theta_2(x, \alpha) \end{pmatrix},$$

and analogously for $d = 3$. The structure for ∇c given by formula (4) leads to the design of a modular toolbox code with the three following mutually independent modules:

$$\text{Distance module} + \text{Image module} + \text{Transformation module} \quad (5)$$

dealing respectively with the task related to the distance functional, the extension and differentiation of the image, and the application of the transformation and its Jacobian. Briefly, the distance module must provide the evaluation of the image distance and the computation of the gradient of the selected distance. The image module comprises an image extrapolation and interpolation model on which the corresponding image gradient will be processed. The transformation module is devoted to deforming the images and computing the Jacobian of the related deformation. For more details on the construction of this modular image registration algorithm, we refer to [25]. This modular structure facilitates not only the construction of the toolbox code but also its extension and use. Indeed, users can easily extend it by implementing any image distance, as well as their own image differentiation scheme, and choose their favorite class of transformation, by simply locally modifying the corresponding module.

The main contribution to the computational cost of the proposed method comes from solving the nonlinear optimization problem (3). In order to make this step cheaper, a multiscale technique can be applied [3, 39], which speeds up the convergence of the co-registration procedure (3) by starting the optimization on lower resolution versions of the images to find a good enough transformation, and then by restarting the process within higher resolution.

In our implementation, we use the image registration toolbox presented in [25] which is equipped with several options of different image distances (including the well-known distances listed in Section 2.2) and image interpolation models. In the following we provide a description of the choices made within the different modules.

2.1. The transformations class

Different choices for the transformation class have been considered in the literature, including rigid transformations (translations, rotations), affine transformations, as well as more general non-rigid transformations (parametric [38], splines [7], B-splines [37, 39, 41, 42], thin plate splines [8]). Assuming, for simplicity, that $\Omega = [0, 1]^d$, we look here for a non rigid deformation in the class of *interpolating scaling functions*. More precisely, we assume that $\Theta_i(x; \alpha)$, belongs to the d dimensional space obtained by tensor products starting from the restriction to $[0, 1]$ of the one dimensional space $V_j \subset L^2(\mathbb{R})$

$$V_j = \text{span}\{\tau_{j,k}(x), k \in \mathbb{Z}\}, \quad \text{where} \quad \tau_{j,k}(x) = \tau(2^j x - k),$$

obtained by rescaling of factor 2^j the integer translates of the the Deslaurier-Dubuc scaling function τ of order $2L + 1$. We recall that the latter can be obtained as autocorrelation of the Daubechies scaling function of order L (see [11]), that we denote φ_L :

$$\tau(x) = \int_{\mathbb{R}} \varphi_L(y) \varphi_L(y - x) dy.$$

We say that the function τ is interpolating, because it verifies $\tau(0) = 1$, $\tau(n) = 0$ for $n \neq 0$. As a consequence a function w in V_j can be written as $w = \sum_k w(k2^{-j})\tau_{j,k}$. The properties of this class of functions are extensively described in [15] (see also [4, 5]). The most relevant here is that they can be organized as a hierarchical structure. More precisely, it is possible to prove that $V_j \subset V_{j+1}$ (in the wavelet terminology we say that the sequence V_j , $j \in \mathbb{Z}$ is a *multiresolution analysis*), and that V_j can be written as the span of a hierarchical basis, itself consisting in translated and rescaled versions of τ :

$$V_j = \text{span}\{\tau_{j_0,k}, k \in \mathbb{Z}\} \cup_{\ell=j_0+1}^j \{\tau_{\ell,2k+1}, k \in \mathbb{Z}\}. \quad (6)$$

We point out that there is a bijective correspondence between the functions of the hierarchical basis and the points of the dyadic grid $k2^{-j}$, $k \in \mathbb{Z}$ ($\tau_{\ell,n}$ corresponds to the point $n2^{-\ell}$ in which it assumes the value 1). This hierarchical structure makes it easy to define non uniform approximation subspaces by simply undersampling the dyadic grid in such a way that it is finer in the regions that need to be better described, and taking the span of the corresponding basis functions. The space V_j can be adapted to the interval $[0, 1]$ (this is quite technical and we refer to [15] for more details) in such a way to maintain the hierchical structure (6) and the bijective correspondence of the basis functions with the dyadic points $k2^{-j}$, $k = 0, \dots, 2^j$ ($k2^{-j} \in [0, 1]$). We will exploit this property later on by using a transformation space adapted to better describe the boundaries of the atlas mesh.

Note that $\nabla_\alpha \theta(x; \alpha)$ depends on x but is independent of α . This is an immediate consequence of the linearity of the expansion (2) in the parameter and holds for any choice of the basis for the transformation space considered. Then $\nabla_\alpha \theta(x; \alpha)$ can be computed once and for all at the points of the grid matrix \mathcal{G} (see Section 2.3).

2.2. The image distance

The role of the functional d is to measure the discrepancy between the two images. In the literature, there are numerous functionals for this purpose (or the equivalent purpose of measuring the similiarity between the two images). Even though such functionals are usually referred to as *image metrics*, most of them are not metrics in the mathematical sense of the term.

The simplest one is the least square (LS) distance, which is the natural distance induced by the continuous image framework based on L^2 -norm. More precisely, for any two images R, R' defined on the same domain Ω , the LS distance is given by

$$LS(R, R') = \frac{1}{2} \|R - R'\|_{L^2(\Omega)}^2 := \frac{1}{2} \int_{\Omega} |R(x) - R'(x)|^2 dx.$$

This very simple (and cheap) distance functional turns out to be surprisingly effective in the absence of noise (see [26]), and it is the one that we use for most of our tests. In more challenging conditions (noisy data), one can resort to more and more sophisticated (and more computationally expensive) metrics, such as the structural similarity index (SSIM, [9, 45]), which independently measures the similarity of luminance, contrast and structure, the Human visual system distance (HVS) [28] (a Fourier weighted L^2 distance with weight suitably defined in order to mimic the response of the human eye), the wavelet normalized root mean square error (WNRMSE) [1], a generalization of SSIM based on the wavelet transform of the images, and the mutual information [39], an extremely expensive but also extremely robust functional mutated from information theory, that can be employed for inter modal registration, which occurs when needing to compare images acquired by different techniques.

2.3. The image model

We recall that both the reference image R and the template image T , in their digital form, are represented as d -dimensional arrays that we interpret as the average of the continuous images over each pixel/voxel. For simplicity we consider the same mesh size in all directions, which corresponds to a square or cubic image domain and to a square or cubic pixels. Then we consider digital images with a $N \times N$ or $N \times N \times N$ resolution. We let $h = 1/N$ and we set

$$\mathcal{G} = \{x^k = ((k_1 - .5)h, \dots, (k_d - .5)h), k = (k_1, \dots, k_d) \in \{1, \dots, N\}^d\}.$$

We denote by pix_k the corresponding pixels

$$\text{pix}_k = [h(k_1 - 1), hk_1] \times \dots \times [h(k_d - 1), hk_d], k = (k_1, \dots, k_d) \in \{1, \dots, N\}^d.$$

In order to compute the cost function $c(\alpha)$ we need to sample the transformed image $\hat{T} = \mathcal{E}T \circ \theta$, with θ any admissible transformation at the points of the grid \mathcal{G} or to compute its average at the corresponding pixel. Then, we have to evaluate the average of $\mathcal{E}T \circ \theta$ over the pixels. Applying the simplest one point quadrature rule (which is the choice that we made) reduces to approximating the average over a pixel with the value of $T \circ \theta$ at the center x^k of the pixel itself.

We then need to define the extension operator \mathcal{E} , allowing us to compute the values of T outside its domain of definition. Our choice is to use a periodicity condition. Assuming that $\Omega = [0, 1]^d$, we set $\mathcal{E}T(x) = T(x - \lfloor x \rfloor)$. This introduces an artificial step gradient at the boundary of Ω , whenever T does not actually satisfy periodic boundary conditions. First we note that many of the images in medical applications do actually satisfy periodic boundary conditions, since the image edges only see the background that is usually uniform. Secondly, numerical tests on generic images show that with a treatment of the gradient compatible with the periodic boundary conditions (i.e., if the step gradient at the boundary induced by the periodic boundary conditions is indeed correctly computed), this artifact does not hinder the registration algorithms.

We also need to evaluate T at the points $\theta(x^k)$. The simplest way of doing this would be to identify the pixel to which $y_k = \theta(x^k)$ belongs, and use the corresponding value of the digital image. However, this *nearest neighbour approach* results in modeling the image as a piecewise constant, which is not differentiable. Instead, our approach for evaluating T at a generic point $y \in \mathbb{R}^d$ consists of performing a d -cubic interpolation based on a centered stencil, whose points are the centers of the pixels.

The second problem is the computation of the gradient of T at the points of the deformed grid, which we address by performing a second order centered scheme on a refined version of the image obtained by iteratively applying L times the above mentioned d -cubic interpolation.

Remark 2.1. Since we deal with digital images, where high frequency information has been discarded, we can always assume that the original analog image has the needed degree of regularity. However, we need to choose the reconstruction procedure, so that the L^∞ function that we would obtain by reconstructing T at all points of Ω exhibits a sufficient degree of smoothness (C^2 is needed for most optimization algorithms).

2.4. Definition of the patient-specific mesh

Let us now come to the main problem of defining a patient-specific mesh. More precisely, let T be a patient-specific image, defined on Ω , and let $\Omega_0 \subset \Omega$ be the subdomain corresponding to the anatomical structure that we need to mesh, in order to perform some numerical simulation. The standard approach consists in applying some segmentation procedure to T , to single out the subdomain Ω_0 , and then in using a meshing procedure to the latter. If the same problem has to be solved for many different patient-specific geometries, obtained from different images (e.g. when repeating the same simulation on many different patients, or at many different times for the same patient), both the segmentation and the meshing procedures will be repeated several times. The main idea of this paper stems from the observation that, in addition to the segmentation of the image at hand, atlas-based segmentation yields the transformation θ_α (in our case of the form (2)), such that the deformed image $R' = \mathcal{E}T \circ \theta_\alpha$ is (ideally) exactly superposed to the reference image R . The patient-specific anatomical structure Ω_0 is then defined as the deformation of the corresponding template anatomical structure $\widehat{\Omega}_0$ (generally defined as $\widehat{\Omega}_0 = \{x \in \Omega : A(x) = e\}$ with $e \in L$ one of the labels). If we have a mesh $\widehat{\mathcal{M}}_h$ on $\widehat{\Omega}_0$, we easily define a mesh on Ω by simply applying θ_α to the coordinates of the nodes of $\widehat{\mathcal{M}}_h$ and keeping the connectivity as it is, which, once θ_α is computed, makes the actual computation of the new mesh extremely cheap).

Notice that when the patient-specific domain is a non connected domain with different components, corresponding to different anatomical structures, then each component of the atlas mesh can be deformed, either separately or at the same time (depending on the user's goal), by the same transformation θ_α that can be computed only once. From the practical point of view, these sub-steps will be facilitated by recording θ_α for reusing.

2.5. “A priori” vs “a posteriori” adaptivity

An important goal of our work is to reduce, as far as possible, the computational effort involved in the numerical simulation of biophysical phenomena, while preserving the good quality of the solution of the problem at hand, solved on several similar geometries. An important ingredient to reach this goal is the use of adaptivity, and the standard approach would be to apply an adaptive procedure to each patient-specific problem, independently. However, in many situations the different patient-specific problems are qualitatively similar. This leads us to consider an alternative option, namely to build a non uniform “adapted” version of the atlas mesh and transfer it to the patient-specific domain through the transformation θ_α . The idea is to run an adaptive procedure on a “reference” problem, posed on the subdomain $\widehat{\Omega}_0$ derived from the reference image R .

Assuming that we need to solve several instances of patient-specific problems, we represent them in abstract form as

$$\mathcal{A}(u; \Omega_0, \mu) = 0$$

defined on the patient-specific domain $\Omega_0 \subset \Omega$ and possibly depending on a vector μ of patient-specific parameters. The idea is to consider a problem

$$\mathcal{A}(\widehat{u}; \widehat{\Omega}_0, \widehat{\mu}) = 0,$$

where $\widehat{\Omega}_0$ is the reference anatomical structure and $\widehat{\mu}$ is a suitably selected reference value for the parameter vector μ . We can then run an adaptive solver and retrieve the adapted mesh that can be transferred to different patient-specific domains through the transformations θ_α stemming from the atlas segmentation procedure. Figure 3 compares the standard approach (DA) to this new approach, (AD). In many cases (when the singularity of the solution stems from shared geometrical or physiological features), this mesh (which, when considering the patient-specific problem, has been adapted “a priori”), will be enough. In other cases, it will provide a good starting point for a (patient-specific) standard adaptive procedure.

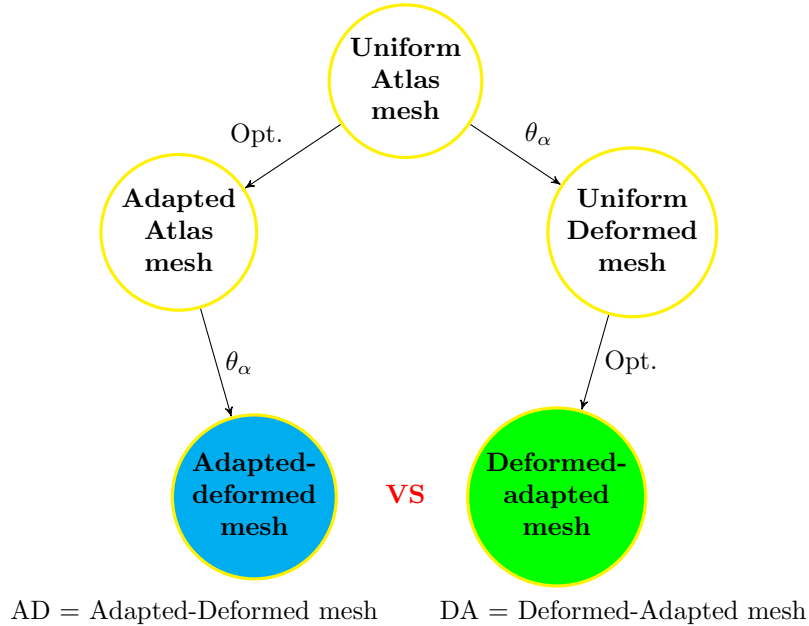


FIGURE 3. What is the optimized scheme for mesh construction ?

3. NUMERICAL TESTS

For all tests in the present paper, we used stock images that we acquired online¹. We implement the atlas construction and co-registration algorithm, using Octave, while the PDE solver for Test 3. is implemented using FreeFem++ [17].

3.1. Test One

First we test our method on a toy example, see Figure 4. We start from the reference image, the x-ray of a knee, and we construct the “patient-specific” knee image T , by manually deforming R in a smooth way, using a mainstream commercial image processing software. We build an atlas on the reference image and we use the transformation obtained from the co-registration between the atlas image and the patient image to deform the atlas meshes to fit the patient-specific domain.

¹©{MonthianRitchan-ad,kravka,andesign1010}@123rf.com




FIGURE 4. Atlas and patient images of a knee. Atlas meshes are deformed to fit the patient-specific domain.

3.2. Test Two

The second test is more challenging, as T and R (displayed in Figure 5) are now two different, independently acquired, stock images of a lateral x-ray of a foot, for which we do not have, a priori, the existence of a transformation that allows the exact superposition of the two. As we can see, besides being in different positions, the two images also display structural differences and the registration procedure is complicated by the presence of very hazy boundaries and by the superposition of different anatomical structures. Focusing on the heel bone, we can see in figure 6a that the registration with uniform transformation space with $j = 2$ based on the LS metric, produces quite poor results. In order to improve the performance, we can refine the transformation space by adding some basis functions at the next level ($j = 3$), corresponding to points close to the boundary of the heel in the Atlas Image (this is accessible, as the Atlas Image is supposed to have been segmented a priori). The result of this approach is shown in Figure 6b. A second possibility, is to use a more refined distance functionals. We chose here to use the WNRMSSE, which has been shown to provide a good compromise between effectiveness and computational cost (see [26]), and show the result in Figure 6c. The two things can be combined, and the WNRMSSE distance can be used in combination with the refined transformation space (see Figure 6d). The result can be further improved by a second level of refinement of the transformation space ($j = 4$).

To select the basis function for the adapted transformation space, we start by taking a coarse uniform mesh (in this case at level $j_0 = 2$) and at each refinement we increase the level by adding, at each new level ℓ the basis functions corresponding to dyadic grid points in a strip ω_Γ^ℓ around the boundary Γ of the atlas anatomical structure, defined by

$$\omega_\Gamma^\ell = \{(x_1, x_2) : \inf_{(\xi_1, \xi_2) \in \Gamma} \max\{|x_1 - \xi_1|, |x_2 - \xi_2|\} \leq 1.5 \cdot 2^{-\ell}\}$$

The results clearly show that both a better choice of the metric function and a local refinement of the transformation space can significantly contribute to improve the quality of the results of the registration step. We would like to remark that the test presented here is indeed quite challenging, as we performed registration of two images of two different individuals, in different position, without any preprocessing. Indeed it is known (see [30]) that the choice and set up of the atlas is of paramount importance for the success of the atlas based segmentation method which underlies our geometry construction and meshing method.

On the negative side, the result also show that simply using the distance between the deformed template and the patient image as cost functional in the optimization procedure may lead to locally losing the shape regularity of the mesh (see Figure 6e).




FIGURE 5. Patient specific meshes




FIGURE 6. Patient specific meshes obtained by the different registration procedures that we tested

3.3. Test three – “A priori” and “a posteriori” adaptivity

In this subsection we test the effectiveness of the “a priori” computation of an adapted mesh, performed accordingly to the idea presented in Section 2.5. We consider the following three problems:

3.3.1. Poisson equation

$$\begin{cases} -\Delta u = 10^{-2} & \text{in } \Omega_0 \\ u = 0 & \text{on } \partial\Omega_0. \end{cases} \quad (7)$$

3.3.2. Elastic equation

$$\begin{cases} -\nabla \cdot \sigma(\mathbf{u}) = \rho \mathbf{g} & \text{in } \Omega_0 \\ \mathbf{u} = \mathbf{0} & \text{on } \partial\Omega_0 \end{cases} \quad (8)$$

which is the equation modeling the equilibrium between efforts and constraints in a given material occupying the computational domain Ω_0 ; in particular, \mathbf{g} is the gravity force term, ρ is the density of the material and $\sigma(\mathbf{u})$ is the constraint tensor associated to the deformation field \mathbf{u} through the Hooke's law

$$\sigma(\mathbf{u}) = \lambda \text{tr}(\varepsilon(\mathbf{u}))I + 2\mu\varepsilon(\mathbf{u}),$$

where λ and μ are the Lamé's coefficients and $\varepsilon(\mathbf{u}) = \frac{1}{2}(\nabla\mathbf{u} + \nabla\mathbf{u}^T)$ is the strain tensor. For our simulation, we set $\rho = 1.9 \times 10^3 \text{ kg m}^{-2}$, $\lambda = 3.81 \times 10^6 \text{ Pa}$ and $\mu = 6.22 \times 10^6 \text{ Pa}$.

3.3.3. Advection-Diffusion-Reaction (ADR) equation

$$\begin{cases} -k\Delta u + \boldsymbol{\beta} \cdot \nabla u + u = f & \text{in } \Omega_0 \\ u = u_{ex} & \text{on } \partial\Omega_0 \end{cases} \quad (9)$$

where $k = 10^{-5}$, $\boldsymbol{\beta} = [0, 1]^T$ and the force term f is chosen such that the exact solution of the problem is $u_{ex}(x, y) = kxy^2$.

These equations are solved on a toy patient-specific domain Ω_0 obtained by the Tibia-fibula image of Figure 4. For the sake of simplicity, the patient-specific domain is supposed to be simply connected. Observe in Figure 7 that the domain is not anatomically correct, as it merges in a single domain the two independent bones. However, as we are not interested in a specific medical application but rather in demonstrating the feasibility of a general meshing strategy, the tests that we are going to perform are still significant. We emphasize that the domain Ω_0 is not trivial since, as we observe in Figure 7, its boundary has positive and negative curvatures that affect the quality of the mesh and thus the mesh where the partial differential equations is solved. We show in Figure 7 the non optimized meshes of the atlas (reference) and patient domains.




FIGURE 7. Non optimized meshes.

All the simulation are computed in FreeFem++, thus the weak solution of the presented problems are looked for by using a finite element linear polynomial space $P1$ for the Galerkin projection. In order to perform the mesh adaptation, both on the reference problem for the “a priori” adaptive mesh computation and on the patient-specific problem, we used the command `adaptmesh` (see [18]), which iteratively generates a sequence of adapted mesh on the solution of a specific problem by starting from a given initial mesh.




FIGURE 8. Adapted meshes and solutions of the Poisson equation.




FIGURE 9. Adapted meshes and solutions of the Elastic equation.

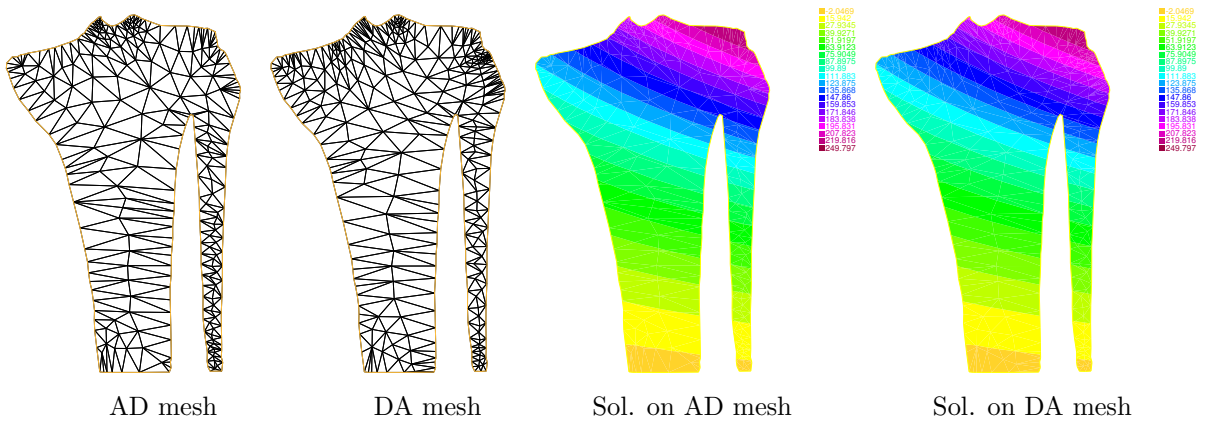


FIGURE 10. Adapted meshes and solutions of the Advection-Diffusion-Reaction equation.

3.3.4. Comparison of the results

In Figures 8, 9 and 10 we compare, for the three model problems, the meshes and the solutions obtained with the two approaches presented in Section 2.5 (resp. AD and DA). From a very qualitative analysis, one cannot notice sensitive differences between the two approaches. We remark that the AD strategy is computationally cheaper than the DA one. As a matter of fact, if in the adapt-then-deform strategy the adaptation is made once for all on the atlas mesh, the deform-then-adapt strategy computes the adaptation of the mesh over the patient specific image for every patient mesh. If the analysis between the final AD and DA meshes underlines that the results obtained on one and on the other mesh is the same, we can conclude by stating that the AD strategy is the better.

We continue the comparison by reporting in Table 1, for the three problems and the two strategies, the attained relative error:

$$\text{Relative error} = \frac{\|u - u^*\|_{L^2(\Omega)}}{\|u^*\|_{L^2(\Omega)}}, \quad (10)$$

where u^* is either the exact solution (for the ADR problem) or the best computed solution (for the other two problems). In particular, when u^* is not the exact solution, it is computed as the quadratic finite element solution on a fine mesh obtained by refining the DA mesh. The results confirm the observation of the previous sections. For Problems 7 and 8 the results obtained by the AD strategy are comparable to those obtained by the (standard) DA strategy, while we see a higher (though not dramatic) loss of accuracy for Problem 9, whose solution is, however, still satisfactorily approximated.

Equation	AD	DA
Poisson	2.10 E-3	1.93 E-3
ADR	4.09 E-4	2.69 E-4
Elastic	2.58 E-3	2.28 E-3

TABLE 1. Relative estimates on adapted-deformed (AD) and deformed-adapted (DA) meshes.

4. CONCLUSION AND PERSPECTIVES

We presented a general strategy for defining the geometry and the mesh for patient-specific numerical simulation of biophysical phenomena. The strategy is based on atlas-based image segmentation and consists in taking advantage of the transformation stemming from the image registration step of atlas-based segmentation, for transferring onto the patient-specific domain, geometry and mesh information pre-defined on the reference (atlas) domain.

The numerical tests performed, give a first confirmation of the feasibility of our strategy, while pointing out some of its limitations, thus suggesting several directions for future research aimed at understanding the field of potential application and at tackling its weaknesses.

Work needs to be put in understanding how to better capture the finest features of the geometry. This issue can be tackled from different sides such as the definition of the distance and the selection of the transformation class. However, it seems important to be able to perform some kind of refinement on the class of transformations. Here, the use of a class of transformation inherently endowed with a multilevel structure (such as the class of interpolating wavelet that we use here) seems quite promising.

Obviously, we would like the deformed mesh to inherit as much as possible the (good) shape regularity properties of the reference mesh. In order to attain this goal, it might be useful to include, in the cost functional, a term penalising the possible lack of smoothness of the transformation. We might for instance think of replacing the functional $c(\alpha)$ with

$$c_\lambda(\alpha) = c(\alpha) + \lambda \|\nabla \theta_\alpha\|_2^2,$$

or, more generally, with

$$\tilde{c}(\alpha) = d(T \circ \theta_\alpha, R) + \rho(\theta_\alpha)$$

where $\rho(\theta)$ is some measure of the distortion resulting from the transformation θ_α .

Finally, we plan to extend the method to the three dimensional case; however, since the computation will require memory and CPU time, we intend to tackle this extension by using a high-performance computational platform.

REFERENCES

- [1] M.G. Albanesi, R. Amadeo, S. Bertoluzza, G. Maggi. *A new class of wavelet-based metrics for image similarity assessment*. Journal of Mathematical Imaging and Vision, Vol. 60, Issue 1, pp 109–127, 2018.
- [2] C. Baillard, P. Hellier and C. Barillot. *Segmentation of brain 3D MR images using level sets and dense registration*, Medical Image Analysis, Vol. 5, No. 3, pp. 185–194, 2001.
- [3] G. Maggi. *Registration algorithms for medical images*. Ph.D Thesis. 2013.
- [4] S. Bertoluzza, G. Naldi. *A wavelet collocation method for the numerical solution of partial differential equations*. Appl. and Comput. Harm. Anal., Vol 3, pp. 1–9, 1996.
- [5] S. Bertoluzza, V. Perrier. *A new construction of boundary interpolating wavelets for fourth order problems*. Acta Appl. Math., Vol152, pp. 33–56, 2017.
- [6] O. V. Besov. *On a family of function spaces. embedding theorems and extensions*. Dokl. Acad. Nauk SSSR, Vol 126, pp. 1163–1165, 1959.
- [7] P. Bézier. *Courbes et surfaces pour la CFAO*, Techniques de l’ingénieur Applications des mathématiques, No A1440, 1992.
- [8] F. L. Bookstein. *Principal warps: Thin-plate splines and the decomposition of deformations*. IEEE Transactions on pattern analysis and machine intelligence, Vol 11, No 6, pp. 567–585, 1989.
- [9] D. Brunet, E. R. Vrscay, Z. and Wang. *On the mathematical properties of the structural similarity index*. IEEE Transactions on Image Processing, Vol 21(4), pp. 1488–1499, 2012.
- [10] D. L. Collins, C. J. Holmes, T. M. Peters and A. C. Evans. *Automatic 3D modelbased neuroanatomical segmentation*, Human Brain Mapping, Vol. 3, No. 3, pp. 190–208, 1995.
- [11] I. Daubechies, Ten Lectures on Wavelets, Society for Industrial and Applied Mathematics Philadelphia, ISBN:0-89871-274-2, 1992.
- [12] B. Delaunay. *Sur la sphre vide. A la mmoire de Georges Voronoï*. Bulletin de l’Académie des Sciences de l’URSS, Classe des sciences mathématiques et naturelles, Issue 6, pp. 793–800, 1934.
- [13] R. A. DeVore and V. A. Popov. *Interpolation of Besov spaces*, Transactions of the American Mathematical Society, Vol 305(1), pp. 397–414, 1988.
- [14] R. A. DeVore and R. C. Sharpley. *Besov spaces on domains in \mathbb{R}^n* . Transactions of the American Mathematical Society, Vol 335(2), pp. 843–864, 1993.
- [15] D. Donoho. *Interpolating wavelet transform*. 1992.
- [16] J. C. Gee, M. Reivich and R. Bajcsy. *Elastically deforming a three-dimensional atlas to match anatomical brain images*, Journal of Computer Assisted Tomography, Vol. 17, No. 2, pp. 225–236, 1993.
- [17] F. Hecht, *New development in FreeFem++*. J. Numer. Math. 20, no. 3-4, 251265. 2012.
- [18] F. Hecht, *A fewsnags in mesh adaptation loops*. In: Hanks B.W. (eds) Proceedings of the 14th International Meshing Roundtable. Springer, Berlin, Heidelberg, 2005.
- [19] C. Johnson. *Numerical Solution of Partial Differential Equations by the Finite Element Method*. Lund, Sweden: Studentlitteratur, 1987.
- [20] C. Johnson and K. Eriksson. *Adaptive Finite Element Methods for Parabolic Problems I: A Linear Model Problem*. SIAM J. Numer. Anal, 28, (1991), pp. 4377.
- [21] H. Kalinić. *Atlas-based image segmentation: A Survey*. Department of Electronic Systems and Information Processing, University of Zagreb. (2008)
- [22] C. T. Kelley. *Iterative Methods for Linear and Nonlinear Equations*. SIAM, Philadelphia, 1995.
- [23] N. Kovacevic, N. J. Lobaugh, M. J. Bronskill, B. Levine, A. Feinstein, and S. E. Black. *A robust method for extraction and automatic segmentation of brain images*, NeuroImage, Vol. 17, No. 3, pp. 1087–1100, 2002.
- [24] P. G. Lemarié. *Some remarks on wavelet theory and interpolation*, Preprint, Mathématiques, Paris XI, Orsay, 1991.
- [25] S. Bertoluzza, G. Maggi and S. Tomatis. *A modular registration algorithm for medical images*. International Conference Image Analysis and Recognition, Springer, pp 467–474, 2013.
- [26] Maggi, G., *Multilevel Registration Algorithms for Medical Images*, PhD Thesis. University of Pavia, 2012
- [27] R. Malladi, J. A. Sethian and B. C. Vemuri. *Shape modeling with front propagation: A level set approach*, IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 17, No. 2, pp. 158–175, 1995.

- [28] J. L. Mannos and D. J. Sakrison. *The effects of a visual fidelity criterion on the encoding of images*. IEEE Transactions on Information Theory, Vol IT-20, pp. 525–536, 1974.
- [29] J. Nocedal and S. Wright. *Numerical Optimization*. Springer-Verlag New York, 2006.
- [30] T. Rohlfing, R. Brandt, R. Menzel, D.B. Russakoff, and C.R. Maurer, Jr., *Quo vadis, atlas-based segmentation* in Handbook of biomedical image analysis, Volume 3: Registration Models. Springer Science & Business Media, 2005.
- [31] A. Sarti, C. O. De Solrzano, S. Locket and R. Malladi. *A geometric model for 3D confocal image analysis*, IEEE Transactions on Biomedical Engineering, Vol. 47, No. 12, pp. 1600–1609, 2000.
- [32] Z. Shi and Z. Bao. *Fast image coding of interval interpolating wavelets*, Wavelet Application IV, SPIE, Vol. 3078, pp. 240–253, 1997.
- [33] I. J. Shoenberg. *Cardinal Interpolation and spline functions: II. Interpolation of data of power growth*, Journ. Approx. Theory, Vol 6, pp. 404–420, 1972.
- [34] J. S. Suri, D. L. Wilson and S. Laxminarayan. *Handbook of biomedical image analysis, Volume 1: Segmentation Models, Part A*. Springer Science & Business Media, 2005.
- [35] J. S. Suri, D. L. Wilson and S. Laxminarayan. *Handbook of biomedical image analysis, Volume 2: Segmentation Models, Part B*. Springer Science & Business Media, 2005.
- [36] J. S. Suri, D. L. Wilson and S. Laxminarayan. *Handbook of biomedical image analysis, Volume 3: Registration Models*. Springer Science & Business Media, 2005.
- [37] P. Thévenaz, T. Blu, and M. Unser. *Interpolation revisited*. IEEE Transactions on Medical Imaging, 19(7), pp. 739–758, 2000.
- [38] D. A. W. Thompson. *On Growth and Form*. Cambridge University Press. 1917.
- [39] P. Thévenaz and M. Unser. *Optimization of mutual information for multiresolution image registration*. IEEE Transactions on Image Processing, 9(12), pp. 2083–2099, 2000.
- [40] A. Tsai, Jr. A. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, W. E. Grimson and A. Willsky. *A shape-based approach to the segmentation of medical imagery using level sets*, IEEE Transactions on Medical Imaging, Vol. 22, No. 2, pp. 137–154, 2003.
- [41] M. Unser, A. Aldroubi, and M. Eden. *B-spline signal processing: Part I – theory*. IEEE Transactions on Signal Processing, 41(2), pp. 821–833, 1993.
- [42] M. Unser, A. Aldroubi, and M. Eden. *B-spline signal processing: Part II – efficient design and applications*. IEEE Transactions on Signal Processing, 41(2), pp. 834–848, 1993.
- [43] M. W. Vannier, T. K. Pilgram, C. M. Speidel, L. R. Neumann, D. L. Rickman and L. D. Schertz. *Validation of magnetic resonance imaging (MRI) multispectral tissue classification*, Computerized Medical Imaging and Graphics, Vol. 15, No. 4, pp. 217–223, 1991.
- [44] P. Viola and W. M. Wells III. *Alignment by maximization of mutual information*. International journal of computer vision, Vol 24(2), pp. 137–154, 1997.
- [45] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. *Image quality assessment: from error visibility to structural similarity*. IEEE Transactions on Image Processing, Vol 13 (4), PP. 600–612, 2004.
- [46] W. M. Wells, W. E. L. Grimson, R. Kikinis and F. A. Jolesz. *Adaptive segmentation of MRI data*, IEEE Transactions on Medical Imaging, Vol. 15, No. 4, pp. 429–442, 1996.
- [47] Jr., A. Yezzi, S. Kichenassamy, A. Kumar, P. Olver and A. Tannenbaum. *A geometric snake model for segmentation of medical imagery*, IEEE Transactions on Medical Imaging, Vol. 16, No. 2, pp. 199–209, 1997.