

Geometry Driven Volumetric Registration

Gheorghe Postelnicu, Lilla Zollei, Rahul Desikan, and Bruce Fischl

MGH/MIT/HMS Athinoula A. Martinos Center for Biomedical Imaging,
Charlestown, MA

{postelni,lzollei,rahul,fischl}@nmr.mgh.harvard.edu

Abstract. In this paper, we propose a novel method for the registration of volumetric images of the brain that attempts to maximize the overlap of cortical folds. In order to achieve this, relevant geometrical information is extracted from a surface-based morph and is diffused throughout the volume using the Navier operator of elasticity. The result is a volumetric warp that aligns the folding patterns.

1 Introduction

Pairwise brain registration is one of the active areas of research in the medical imaging community. Different algorithms have tackled the generic problem of registering information from two brain scans in various ways. Volumetric registration (see [1] for a survey) seeks a 3D deformation field which is driven by either raw intensity information or features derived from image intensities. A different approach is to extract geometric features from surface models of structures such as the neocortex, and to reformulate the complex correspondence problem in a *surface matching* framework.

Each of these approaches has advantages and weaknesses. Surface-based methods [2,3,4] have been shown to accurately align the highly complex folding pattern of the human cerebral cortex, and to result in increased statistical power presumably due to their alignment of functionally homologous regions across subjects. This accuracy stems from the direct use of geometric information that is generally unavailable to volumetric methods and the relatively close relationship between folding patterns and functional properties of the neocortex. Conversely, volumetric methods [5,6,7,8], while frequently failing to align corresponding cortical folds, provide a correspondence field in the whole brain, and align subcortical and ventricular structures as well as the cortex (regions that are outside of the domain of classical surface-based registrations).

In this paper, we propose a method which combines the two approaches. This is done by integrating surface-based information into a volumetric registration procedure. The result is a 3D deformation field which aligns the folding patterns of the the two scans. While the idea of using surface registration to drive volumetric deformation fields is not new [9,10,11], this is the first paper to explicitly concentrate on the accurate registration of the cortical sheet in 3D space.

The current work can be seen as growing out of the non-linear registration literature that aims at integrating prescribed displacements into a volumetric

morph: [9] proposes a linear and incremental method for performing volumetric deformations, and [10,11,12] use various types of surfaces and matching algorithms that are then interpolated in the rest of the image to yield a dense deformation field, while [13] uses explicitly extracted sulcal traces together with feature vectors. Perhaps the closest related work to this article is by Liu et al. [14], where the geometrical information carried by the brain surfaces is explicitly used in the registration process. However, it should be mentioned that the surface similarity is maximized after the volumetric warping, which we believe makes it susceptible to local minima.

2 Methods

We present the process of registering two structural brain scans, *fixed* and *moving*. Each of the scans is independently processed to obtain an accurate topologically correct reconstruction of the cortical surfaces (see [15,16,17,18,19] for details). Then, we perform a surface-based registration, independently for each of the surfaces (4/brain - left and right pial and gray/white respectively). This registration takes place in spherical coordinates and aims at maximizing the similarity of the folding patterns, while remaining topologically correct and controlling the amount of allowed metric distortions.

The result of the surface registration algorithm provides the input to the volumetric registration we present in this paper. Using a regularizer from the theory of elasticity, we build a displacement field which is driven by the surface registration. The resulting morph yields surfaces that are as close as possible to the target surfaces, while being topologically correct and respecting the anatomical variability between individuals.

2.1 Surface Registration

The surface-based registration of our choice is briefly described below [3]. The first step is to transform the cortical surface into a spherical representation with moderate metric distortions [20]. The algorithm aims at iteratively minimizing the following energy in the spherical space:

$$J = J_p + \lambda_A J_A + \lambda_d J_d \quad (1)$$

where J_p measures the alignment, based on the cortical depth and the curvature information, while the other two terms act as regularizers. J_A is a topology-preservation term, while J_d controls the amount of metric distortion allowed:

$$J_A = \frac{1}{2T} \sum_{i=1}^T (A_i^n - A_i^0)^2 \quad J_d = \frac{1}{4V} \sum_{i=1}^V \sum_{n \in N(i)} (d_{in}^n - d_{in}^0)^2 \quad d_{in} = \|x_i - x_n\|$$

Here, superscripts denote time with 0 being the starting point, T and V are the number of triangles and vertices in the tessellation, x_i^n is the position of

vertex i at iteration n , $N(i)$ is the set of neighbors of the i^{th} vertex and A_i denotes the *oriented* area of triangle i . The result of the surface registration is a 1-to-1 mapping that transports each surface of the fixed scan to its counterpart surface in the moving image $F_{reg} : \mathcal{S}_{fixed} \rightarrow \mathcal{S}_{moving}$, where S can be any of the left/right pial/white surfaces of the brain.

2.2 Volumetric Warping

When the surface registration is completed, we obtain a displacement vector field which provides a 1-to-1 mapping between the hemisphere surfaces of the fixed and moving brain scans in the Euclidean space. We now show how to diffuse this vector field from the cortical surfaces to the rest of the volume.

Let Ω be the source image domain. We define an arbitrary transformation $\phi : \Omega \rightarrow \mathbb{R}^3$ of the source image as: $\phi(x) = x + u(x)$ where $u : \Omega \rightarrow \mathbb{R}^3$ denotes the displacement field. The goal here is to find a function ϕ such that $\phi(x_{surf}^{fixed}) = F_{reg}(x_{surf}^{fixed})$, for any $x_{surf}^{fixed} \in \mathcal{S}$, where \mathcal{S} is one of the surfaces of the fixed brain scan. Since the surfaces represent a space of co-dimension 1, in order for this problem to be well-posed, we impose an additional regularity constraint. We require that the displacement field we are searching to be an *elastic deformation*, i.e. a smooth, orientation-preserving deformation which satisfies the equations of static equilibrium in elastic materials. This means that we (additionally) require u to satisfy

$$\mathcal{L}(u) = 0 \tag{2}$$

where \mathcal{L} is an operator we define below.

The choice of the operator \mathcal{L} and the discretization method to numerically extrapolate the displacement field has numerous solutions that have been proposed in the registration literature. Some of the better known are the thin plate splines, proposed by Bookstein [21], or the free-form deformations proposed by Rueckert et al [22]. We have chosen to use the Navier operator from the linearized elasticity theory together with the finite element method. This choice was motivated by the high level of flexibility needed in order to satisfy the constraints imposed by the displacement fields obtained on the 4 surfaces of the brain image.

2.3 Elasticity Operator

In order to solve the problem stated above, we use the equilibrium equation for elastic materials. This states that at equilibrium, the elastic energy equals the external forces applied to the body $\mathcal{L}(u) = f$ [23].

$$\mathcal{L}(u) = -\text{div} \left[(I + \nabla u) \hat{S} \right] \tag{3}$$

where the second Piola-Kirchoff stress tensor $\hat{S} : \bar{\Omega} \rightarrow \mathcal{M}^3$ is defined as

$$\hat{S} = \lambda \text{tr}(\hat{E})I + 2\mu \hat{E} \quad \text{and} \quad \hat{E} = \frac{1}{2} (\nabla u^T + \nabla u + \nabla u^T \nabla u)$$

is the Green-St. Venant strain tensor. Here λ and μ are the Lamé elastic constants that characterize the elastic properties of an isotropic material. The linear approximation to the above operator uses the Fréchet derivative of \mathcal{L}

$$\mathcal{L}(u) = \mathcal{L}(0) + \mathcal{L}'(0)u + o(u) \Rightarrow f = \mathcal{L}(u) \approx \mathcal{L}'(0)u$$

since no deformation occurs in the absence of external forces. Finally, $\mathcal{L}'(0)u$ is computed by dropping the non-linear terms in $\mathcal{L}(u)$, which results in $\mathcal{L}'(0)u = -\operatorname{div} S$ with $S = \lambda \operatorname{tr}(E)I + 2\mu E$ the linearized stress tensor and $E = \frac{1}{2}(\nabla u + \nabla u^T)$ the linearized strain tensor. Hence, the linear approximation of (2) can be written as

$$\mathcal{L}'(u) = -\operatorname{div} S = f \tag{4}$$

The main drawback of (4) is that it is only valid for small-magnitude deformations. To overcome this, we implement an extension of the linear model, as presented in [9]. Namely, given external forces that describe large displacements, one can iteratively solve for small linear increments using the linearized Navier equation (4): $\mathcal{L}(u^{n+1}) = \mathcal{L}'(u^n)(u^{n+1} - u^n) + o(u^{n+1} - u^n)$ or, by neglecting the last term,

$$f^{n+1} - f^n = \mathcal{L}(u^{n+1}) - \mathcal{L}(u^n) \approx \mathcal{L}'(u^n)(u^{n+1} - u^n) \tag{5}$$

Using this iterative process, the solution of (5) converges to the solution of (2) (see [23] for the proof).

The Lamé constants λ and μ are specified as functions of *Young's modulus of elasticity* \mathcal{E} and the *Poisson ratio* ν : $\lambda = \frac{\mathcal{E}\nu}{(1+\nu)(1-2\nu)}$ and $\mu = \frac{\mathcal{E}}{2(1+\nu)}$. In all our experiments we used $\mathcal{E} \equiv 1$ and $\nu = 0.3$.

2.4 Finite Element Method

Using the notation introduced in the previous subsection, equation (4) can be re-written as a minimization problem by considering the potential energy of an elastic body submitted to externally applied forces

$$\mathbb{E} = \frac{1}{2} \int_{\Omega} \left[\lambda \|\nabla \cdot u\|^2 + 2\mu \sum_{i,j=1}^3 e_{ij}^2(u) \right] dx + \int_{\Omega} f \cdot u dx \tag{6}$$

with $e_{ij}(u) = \frac{1}{2}(\partial_i u_j + \partial_j u_i)$.

One common way to tackle a minimization problem for (6) is to use the Finite Element Method. This consists in dividing the domain Ω using an assemblage of discrete finite elements interconnected at nodal points at the element boundaries (tetrahedra in our case). The continuous field u within each element can be approximated as a linear function of the displacements at the nodal points - the computation is detailed in [24]. The first term in (6) results in a sparse symmetric matrix K , also called the stiffness matrix, whereas the external forces result in a vector F . By using the linear approximations mentioned, (6) can be re-written

$$\mathbb{E}_{mesh} = U^T K U + U^T F \tag{7}$$

hence by minimizing (6) wrt. the nodal displacements $\frac{dE_{mesh}(U)}{dU_i} = 0$, equation (4) can be written as

$$KU = F \quad (8)$$

2.5 Specifying the External Forces

All that remains to be specified for equation (8) is the vector of external forces F . As mentioned previously, the input to our problem is actually a set of prescribed displacements.

One way of implementing prescribed displacements as external forces when solving a FEM problem is described in [9]. It consists of modifying the matrix K and the vector F so that the value of the variable U_i we want to prescribe is forced. This is implemented by setting $K_{ij} = \delta_{ij}$ and $K_{ji} = \delta_{ji} \forall j$ and subtracting the appropriate quantity on the right hand side. However, numerical experiments showed that this way of constraining the stiffness matrix is too strong and it can cause topology problems (i.e. noninvertible regions). Indeed, in the current problem domain, where anatomical differences between the two brains are to be expected (such as a split fold), such hard constraints are undesirable, as they result in overfitting the warp field.

We have instead opted for an implementation of prescribed displacements using penalty weighting. This means that we use the displacements given by the surface registration F_{reg} to modify equation (4). Indeed, without external constraints, the elasticity problem simply reads $KU = 0$, with a trivial solution.

Suppose we want to impose a constraint $u(x) = v$. We start by determining the tetrahedron T_i such that $x \in T_i$ and we know that we must have $\sum_{j=1}^4 N_j^{T_i}(x)U_j^{T_i} = v$, where N_j are the barycentric coordinates of x in T_i . Then we impose the condition by modifying (7) so that the energy becomes

$$E_{mesh}^{mod} = U^T K U + \alpha U^T (A_i U - B_i) \quad (9)$$

where A_i are the barycentric coordinates of $x \in T_i$ and $B_i = v$ (which are immersed in the global array U using the correspondence table of the nodes of tetrahedron T_i). In (9), α is the weight placed on the constraints. In all of our experiments we used a constant weight for all displacement vectors, although this could change in the future. For instance, a surface similarity measure could be used to better condition the morph.

2.6 Handling Topology Problems

After each iteration of the elastic solver, we check for potential topology problems, i.e. tetrahedra with negative Jacobian. Indeed, in spite of using the incremental model and solving the elasticity equation in more iterations, topology problems do appear. There are two general techniques used for performing mesh untangling: smoothing or, more generally, solving an optimization problem aiming at locally fixing tangles (see [25] for an algorithm that combines the two).

In our case, the goal pursued was to solve the tangles from the deformed mesh (i.e. inverted tetrahedra in the deformed configuration) while modifying the displacement field as little as possible. This task is facilitated by the fact that in the absence of displacements, the initial mesh is guaranteed to be topologically correct. Thus, we opt for a local smoothing-based algorithm, which we formulate as the solution of a local linear elastic problem.

We segment the regions with topological defects into local connected components. We then solve a linear elastic equation within each local cluster, the external forces being the displacements on the boundary of the cluster. Finally, we discard the initial displacements inside the cluster and replace them with the solution of the system.

We repeat the above clustering procedure until there are no more tetrahedra with topology problems. Although there is no formal proof of it, the above procedure has solved all the topological problems we were experiencing in the tests we ran thus far.

2.7 Implementation

To resume, the execution of the pipeline described in this section results in the following:

1. compute surfaces for each of the brain images (2 surfaces per hemisphere - pial and gray/white);
2. perform surface registration for each of the surfaces independently in spherical coordinates;
3. recover sparse displacement fields $x_i^{fixed} \rightarrow y_i^{moving}$;
4. regress out affine transform from the displacement field A ; this results in the updated sparse displacement field $x_i^{fixed} \rightarrow A^{-1} \left(y_i^{moving} \right) = z_i$;
5. apply linear incremental model in n steps; i.e. loop $j = 1 \dots n$
 - (a) get current morphed positions $\phi_j \circ \phi_{j-1} \circ \dots \circ \phi_1(x_i) = x_i^j$ and create sparse displacement field $v_i^j = \frac{1}{n-j+1}(z_i - x_i^j)$
 - (b) create tetrahedral mesh based on current surface positions and initialize the stiffness matrix and the external forces;
 - (c) solve the linear system;
 - (d) handle potential topology problems.

We use TetGen [26,27] to build a Delaunay tetrahedral mesh which is adapted to the input surfaces and PETSc [28,29] to solve the linear system at each step. Generally, we place a constraint on the volume of the tetrahedra that are near the surfaces (so that they have a lower volume), since it is expected that the morph will require greater flexibility in those areas.

It should also be noticed that prior to the elastic registration we apply an affine registration. This is motivated by the fact that the linearized version of the elastic operator does not satisfy the axiom of material frame-indifference [23,30]. This means that the linearized elastic energy increases when the object is rotated. This occurs because the linearized operator drops the quadratic terms. Even though

the linear incremental model approximates the Navier operator as the number of steps $n \rightarrow \infty$, applying the affine transform prior to performing the elastic registration allows to decrease n , which results in a significant computational time gain.

3 Results

We illustrate the utility of the method through two experiments. First, we show how the elastic warping can successfully recover the nonlinear deformation caused in the brain due to autopsy and fixation. Subsequently, we present a more ambitious experiment, where we use the pairwise surface correspondence to obtain a warp between brain images of different subjects. As will be seen, this is a more challenging problem, being inherently ill-posed.

3.1 Registration of ex-vivo Scans with the Corresponding in-vivo Scan

In this case, the imaging protocol for the ex-vivo tissue is different due to the reduced T1 contrast observed post-mortem, so a Multi Echo Flash protocol is used. This makes the pre-processing required to obtain the surfaces for the ex-vivo images a little more challenging, but does not affect the registration algorithm proposed here.

We present in figure 1 the result of the volumetric warp applied to the ex-vivo image so that it matches the in-vivo one. The resulting correspondence is almost perfect, since the underlying anatomy is the same and the deformation is a truly mechanical one. However, we remind the reader that the correspondence is not perfect near the lateral ventricles, because none of the cortical surfaces we have used crosses that area.

Also, it should be noted that in this case only the surfaces from one hemisphere were used in the process, as we only had one hemisphere available for the ex-vivo imaging. Nevertheless, the resulting match is excellent, and highlights the insensitivity of the procedure to the underlying image contrast, as the geometric features are of course invariant to the contrast properties.

3.2 Inter-subject Registration

In this section, we show results from a comparison of our morphing method with the morph produced by the publicly available version of HAMMER [31]¹. We selected eleven subjects for which we had labels that had been manually drawn on the surface [33] and we morphed each of them on a randomly chosen subject

¹ We would like to mention that the HAMMER version we have uses the gray/white matter segmentation produced by FAST [32] exclusively to produce the attribute vectors. As such, it is possible that results improve with different inputs to the attribute vector.




Fig. 1. Results of the surface-driven morph between an ex-vivo hemisphere and an in-vivo scan of the same subject. Surfaces are from the in-vivo data (pial surface in red and gray/white surface in yellow).

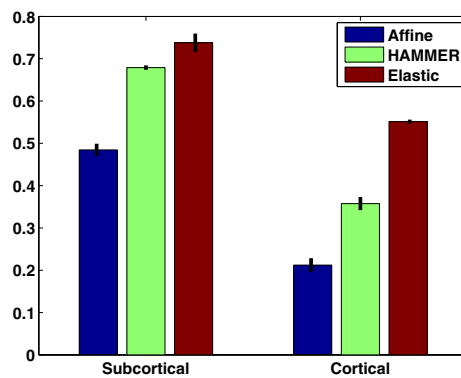


Fig. 2. Average DICE measures with standard error for cortical and subcortical areas over 10 brains. As expected, the measure for cortical areas shows the surface-based morph out-performs HAMMER by 20%. Surprisingly, HAMMER is also out-performed for sub-cortical areas, although our morph does not use any information from these regions.

which played the role of template. We then performed a DICE measure separately for the cortical and subcortical areas and compared the results produced by FLIRT [34] (i.e. 12 DOFs) and HAMMER.




Fig. 3. Visual Comparison of inter-subject registration. Surfaces are from the atlas (pial surfaces in red and gray/white surfaces in yellow). Upper row - FLIRT. Middle row - HAMMER. Bottom row - result of the surface-based morph. It is apparent the HAMMER result is trapped in local minima in certain regions, which does not happen for the surface-based morph. Gyri where large differences occur emphasized with an arrow. It seems reasonable to believe local mis-registrations occur due to a poor initialization.

The validation measures we used were DICE measures on two sets of labels. Using [35], we created a label volume independently for each brain image with each of the major subcortical structures. In addition, we filled the cortical ribbon with manual surface-based labels from [33] for each of the subjects to generate volumetric labels for surface folding patterns. The DICE results were divided into two sets to illustrate the differences in these structures: the sparse displacement fields were explicitly generated to align cortical folds. The degree to which the volumetric morph generated from the surface registration also aligns subcortical regions thus reflects how well predicted the position and shape of subcortical structures are from the folding patterns, with potential implications for neurodevelopment.

To be more specific, the measure we used to compare the degree of overlap between two volumes V_1 and V_2 for a set of labels $S = \{s_i, i = 1 \dots n\}$ is given by

$$\text{DICE}_S(V_1, V_2) = \frac{\sum_{i=1}^n |[V_1 = s_i] \cap [V_2 = s_i]|}{\sum_{i=1}^n |[V_1 = s_i] \cup [V_2 = s_i]|} \quad (10)$$

The morph was executed with $n = 13$ iterations for the linear incremental model and the results are summarized in figure 2. An example of the alignment achieved is shown in figure 3. As was to be expected, the cortical measure is higher by 20% for the surface driven morph than HAMMER (as can be seen in figure 3, HAMMER can get caught in local minima, resulting in over-deformed gyri). Surprisingly, the results also show we outperform HAMMER for subcortical regions, despite the lack of any specific information from these regions in the morph.

4 Conclusion and Future Work

We presented a technique for computing a dense volumetric registration field that was shown to align cortical folding patterns as well as deep brain structures. This was achieved by using a surface-driven morph, together with a regularizer taken from the theory of elasticity to compute a volumetric registration.

Another finding from the present study that has potentially interesting neuroscientific applications is that the alignment of cortical folds also appears to align subcortical structures, indicating that the folds are good predictors of the position/shape of deep brain regions.

In future work, we will use this volumetric field to initialize a volumetric registration algorithm to further align non-cortical structures. We anticipate that this technique will resolve one of the main difficulties with volumetric registration: they do not in general align cortical folding patterns.

Acknowledgments

Support for this research was provided in part by the National Center for Research Resources (P41-RR14075, R01 RR16594-01A1 and the NCRR BIRN Morphometric Project BIRN002, U24 RR021382), the National Institute for Biomedical Imaging and Bioengineering (R01 EB001550), the National Institute for

Neurological Disorders and Stroke (R01 NS052585-01) as well as the Mental Illness and Neuroscience Discovery (MIND) Institute, and is part of the National Alliance for Medical Image Computing (NAMIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 EB005149. Information on the National Centers for Biomedical Computing can be obtained from <http://nihroadmap.nih.gov/bioinformatics>.

References

1. Maintz, J., Viergever, M.: A survey of medical image registration. *Medical Image Analysis* 2(1), 1–36 (1998)
2. Thompson, P., Hayashi, K., et al.: Detecting dynamic and genetic effects on brain structure using high-dimensional cortical pattern matching. In: *Proc. International Symposium on Biomedical Imaging 2002* (2002)
3. Fischl, B., Sereno, M.I., Tootell, R., Dale, A.M.: High-resolution inter-subject averaging and a coordinate system for the cortical surface. *Human Brain Mapping* (8), 272–284 (1999)
4. Davatzikos, C., Bryan, N.: Using a deformable surface model to obtain a shape representation of the cortex. *IEEE Trans. Med. Imag.* 15(6), 785–795 (1996)
5. Bajcsy, R., Kovacic, S.: Multi-resolution elastic matching. *Computer Vision, Graphics, and Image Processing* 46, 1–21 (1989)
6. Davatzikos, C.: Spatial transformation and registration of brain images using elastically deformable models. *Computer Vision and Image Understanding* 66(2), 207–222 (1997)
7. Christensen, G., Joshi, S., Miller, M.: Volumetric transformation of brain anatomy. *IEEE Trans. Med. Imaging* 16(6), 864–877 (1997)
8. Fischl, B., Salat, D.H., van der Kouwe, A.J., Segonne, F., Dale, A.M.: Sequence-independent segmentation of magnetic resonance images. *NeuroImage* 23(1), 69–84 (2004)
9. Peckar, W., Schnorr, C., Rohr, K., Stiehl, H.S.: Parameter-free elastic deformation approach for 2d and 3d registration using prescribed displacements. *J. Math. Imaging Vis.* 10(2), 143–162 (1999)
10. Ferrant, M., Nabavi, A., Macq, B.M., Jolesz, F.A., Kikinis, R., Warfield, S.K.: Registration of 3d intraoperative mr images of the brain using a finite element biomechanical model. *IEEE Trans. Med. Imaging* 20(12), 1384–1397 (2001)
11. Toga, P.M., Toga, A.W.: A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans. Med. Imaging* 15(4), 402–417 (1996)
12. Peckar, W., Schnorr, C., Rohr, K., Stiehl, H.S.: Two-step parameter-free elastic image registration with prescribed point displacements. In: *ICIAP* (1), pp. 527–534 (1997)
13. Collins, D., LeGoualher, G., Caramanos, Z., Evans, A., Barillot, C.: Cortical constraints for non-linear cortical registration. In: *Proceedings of the International Conference on Visualization in Biomedical Computing*, pp. 307–316 (1996)
14. Liu, T., Shen, D., Davatzikos, C.: Deformable registration of cortical surfaces via hybrid volumetric and surface warping. *NeuroImage* 22, 1790–1801 (2004)
15. Segonne, F., Dale, A.M., et al.: A hybrid approach to the skull-stripping problem in mri. *NeuroImage* 22, 1160–1175 (2004)
16. Dale, A.M., Fischl, B., Sereno, M.I.: Cortical surface-based analysis i: Segmentation and surface reconstruction. *NeuroImage* 9(2), 179–194 (1999)

17. Fischl, B., Sereno, M.I., Dale, A.M.: Cortical surface-based analysis ii: Inflation, flattening, and a surface-based coordinate system. *NeuroImage* 9(2), 195–207 (1999)
18. Fischl, B., Liu, A., Dale, A.M.: Automated manifold surgery: Constructing geometrically accurate and topologically correct models of the human cerebral cortex. *IEEE Trans. Med. Imaging* 20(1), 70–80 (2001)
19. Segonne, F., Grimson, E., Fischl, B.: Genetic algorithm for the topology correction of cortical surfaces. *IPMI*, 393–405 (2005)
20. Fischl, B., Dale, A.M., Sereno, M.I., Tootell, R., Rosen, B.: A coordinate system for the cortical surface. *NeuroImage*, (7) (1998)
21. Bookstein, F.: Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11(6), 567–585 (1989)
22. Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., Hawkes, D.: Non-rigid registration using free-form deformations: Applications to breast mr images. *IEEE Trans. Med. Imag.* 18(8), 712–721 (1999)
23. Ciarlet, P.: *Mathematical Elasticity. Three-Dimensional Elasticity. vol. I*, North-Holland, Amsterdam (1988)
24. Zienkewicz, O., Taylor, R.: *The Finite Element Method*. McGraw Hill Book Co. New York (1987)
25. Freitag, L., Plassman, P.: Local optimization-based untangling algorithms for quadrilateral meshes (2001)
26. Si, H., Gaertner, K.: Meshing piecewise linear complexes by constrained delaunay tetrahedralizations. In: *Proceedings of the 14th International Meshing Roundtable 2005*, pp. 147–163 (2005)
27. Si, H.: On refinement of constrained delaunay tetrahedralizations. In: *Proceedings of the 15th International Meshing Roundtable* (2006)
28. Balay, S., Buschelman, K., Eijkhout, V., Gropp, W.D., Kaushik, D., Knepley, M.G., McInnes, L.C., Smith, B.F., Zhang, H.: *PETSc users manual*. Technical Report ANL-95/11 - Revision 2.1.5, Argonne National Laboratory (2004)
29. Balay, S., Gropp, W.D., McInnes, L.C., Smith, B.F.: Efficient management of parallelism in object oriented numerical software libraries. In: Arge, E., Bruaset, A.M., Langtangen, H.P. (eds.) *Modern Software Tools in Scientific Computing*, pp. 163–202. Birkhäuser Press (1997)
30. Picinbono, G., Delingette, H., Ayache, N.: Non-linear and anisotropic elastic soft tissue models for medical simulation. In: *ICRA2001: IEEE International Conference Robotics and Automation*, Seoul Korea (2001) Best conference paper award.
31. Shen, D., Davatzikos, C.: Hammer: Hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. Med. Imaging* 21(11), 1421–1439 (2002)
32. Zhang, Y., Brady, M., Smith, S.: Segmentation of brain mr images through a hidden markov random field model and the expectation maximization algorithm. *IEEE Trans. Med. Imag.* 20(1), 45–57 (2001)
33. Desikan, R., Segonne, F., Fischl, B., et al.: An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. *NeuroImage* 31(3), 968–980 (2006)
34. Jenkinson, M., Bannister, P., Brady, J., Smith, S.: Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* 17(2), 825–841 (2002)
35. Fischl, B.: Automatically parcellating the human cerebral cortex. *Cerebral Cortex* 14, 11–22 (2004)