

Topology Preserving Atlas Construction from Shape Data without Correspondence using Sparse Parameters

Stanley Durrleman[†], Marcel Prastawa^{*}, Julie R. Korenberg^{*}, Sarang Joshi^{*},
Alain Trounev[‡], Guido Gerig^{*,*}

[†] INRIA / ICM, Pitié Salpêtrière Hospital, Paris, France

^{*} SCI Institute, University of Utah, Salt Lake City, USA

^{*} Brain Institute, University of Utah, Salt Lake City, USA

[‡] CMLA - Ecole Normale Supérieure de Cachan, Cachan, France

Abstract. Statistical analysis of shapes, performed by constructing an atlas composed of an average model of shapes within a population and associated deformation maps, is a fundamental aspect of medical imaging studies. Usual methods for constructing a shape atlas require point correspondences across subjects, which are difficult in practice. By contrast, methods based on currents do not require correspondence. However, existing atlas construction methods using currents suffer from two limitations. First, the template current is not in the form of a topologically correct mesh, which makes direct analysis on shapes difficult. Second, the deformations are parametrized by vectors at the same location as the normals of the template current which often provides a parametrization that is more dense than required. In this paper, we propose a novel method for constructing shape atlases using currents where topology of the template is preserved and deformation parameters are optimized independently of the shape parameters. We use an L^1 -type prior that enables us to adaptively compute sparse and low dimensional parameterization of deformations. We show an application of our method for comparing anatomical shapes of patients with Down's syndrome and healthy controls, where the sparse parametrization of diffeomorphisms decreases the parameter dimension by one order of magnitude.

1 Introduction

Shape statistics is fundamental to medical imaging studies, for instance to characterize normal versus pathological structures or to highlight the effect of treatments on anatomical structures. Usual methods build an average shape model, called template, that is representative of a shape ensemble, and correspondences between the template and each subject's shape, altogether called an atlas. Differences appear in the shape model used and the way shapes are put into correspondence. Parametric models, such as medial axis representation [12], extract low-dimensional shape features, which are averaged and compared across subjects. Non parametric methods are often based on point correspondences across

shapes with an either fixed or optimized parameterization of the curves or surfaces [2, 7, 9, 13]. Homologous point positions are averaged and 3D-deformations map the average model to the point configuration of each subject, such maps being used for deriving shape statistics [14, 11]. To alleviate the problem of finding point correspondences across shapes, which is not always possible in practice, one can compare shapes using the metric on currents [15]. In this framework, shape statistics are represented by a template current and diffeomorphic 3D mappings that map the template current to each subject’s shape.

Existing atlas construction approaches using the metric on currents [6, 4, 3] suffer from two main limitations. First, the template current is either the superimposition of warped surfaces [6] or a set of disconnected normals or tangents (Dirac currents) [4, 3] and therefore is not given as a mesh. It has never been shown that such templates could be the approximation of a continuous surface. In this sense, these approaches do not preserve the topology of the shapes, namely the continuity of the surfaces and the number of connected components of the shapes in the population, thus limiting its practical use for segmentation purposes, for instance. Second, the template-to-surface deformations are parameterized by momentum or speed vectors that are at the same location as the normal or tangent of the template current. However, there is no reason to link the parameterization of the template shape with the parameterization of the deformations. The most variable parts of the shape are not necessarily the parts that require the finest sampling, e.g. the most folded ones. Therefore, the template shape and the deformations should have independent parameterization.

In this paper, we propose to build a shape atlas using the metric on currents, but in a way which allow the user to fix the topology of the template mesh, namely the number of vertices and the edges connecting them, and to optimize their position in the 3D space. This will be possible due to an optimal control formulation of the diffeomorphic matching problem in the Large Deformation Diffeomorphic Metric Mapping paradigm [10, 5]. In this formulation, the three variables to optimize, namely the template shape, the control points positions and the momentum vectors attached to them that parameterize the deformations, are seen as the initial conditions of a dynamical system. We derive here the analytic expression of the gradient with respect to these variables. Additionally, we use a L^1 prior to select the most relevant subset of control points for the parameterization of the shape variability. Such a sparse and adaptive parameterization will be particularly well suited for statistical purposes. We show an application of our method to a study comparing anatomical features of patients with Down’s syndrome (DS) and healthy controls, where the sparse parametrization of shape variability enable statistical analysis in lower dimensional spaces.

2 Shape Atlas Construction

2.1 Joint Estimation of Template and Deformations

The method aims to estimate a template shape \mathbf{X}_0 and template-to-subjects deformations ϕ_i from a set of shapes $\mathbf{X}_1, \dots, \mathbf{X}_{N_{\text{su}}}$. Each \mathbf{X}_i denotes a vector

containing the positions of the vertices of the input shapes, which may be of different sizes. \mathbf{X}_0 denotes the vertices of the template shape, whose number and connectivity is fixed by the user. Only the positions of the vertices of the template are optimized and not its connectivity, so that the topology of the template is preserved during optimization. In the framework of Fréchet means [11], we estimate one template and N_{su} deformations ϕ_i that minimize the criterion:

$$E(\mathbf{X}_0, \{\phi_1, \dots, \phi_{N_{\text{su}}}\}) = \sum_{i=1}^{N_{\text{su}}} D(\phi_i(\mathbf{X}_0), \mathbf{X}_i) + \text{Reg}(\phi_i), \quad (1)$$

where D denotes the squared distance on currents, $\phi_i(\mathbf{X}_0)$ the deformation of the template shape \mathbf{X}_0 and Reg a measure of regularity of the deformations.

2.2 Parameterization of Deformations

We use a mechanical system of self-interacting particles, called control points, to build dense diffeomorphic deformations [5]. Let $\mathbf{c}_0 = \{c_{0,k}\}$ be a set of control points and $\boldsymbol{\alpha}_0 = \{\alpha_{0,k}\}$ a set of initial momenta of the particles, altogether called the initial state of the system $\mathbf{S}_0 = \{\mathbf{c}_0, \boldsymbol{\alpha}_0\}$. This set of particles moves from time $t = 0$ to $t = 1$ according to following equations of motion:

$$\begin{cases} \dot{c}_k(t) = \sum_{p=1}^N K(c_k(t), c_p(t)) \alpha_p(t) \\ \dot{\alpha}_k(t) = - \sum_{p=1}^N \alpha_k(t)^t \alpha_p(t) \nabla_1 K(c_k(t), c_p(t)) \end{cases}, \quad (2)$$

which are such that the energy of the system $\sum_{i,j} \alpha_i(t)^t K(c_i(t), c_j(t)) \alpha_j(t)$ is conserved in time. The kernel K models the interaction forces among the particles. These equations describe the evolution of the state of the system $\mathbf{S}(t) = \{c_k(t), \alpha_k(t)\}$ and can be written in short: $\dot{\mathbf{S}}(t) = F(\mathbf{S}(t))$ with the initial condition $\mathbf{S}(0) = \mathbf{S}_0$.

The motion of the control points defines a diffeomorphism of the whole 3D space [8]. The speed at position x_0 interpolates the speed of the control points:

$$\dot{x}(t) = \sum_{k=1}^N K(x(t), c_k(t)) \alpha_k(t). \quad (3)$$

This equation shows that the rate of decay of the kernel determines the size of the neighborhood that is “pulled” by each control point. It can be written in short as $\dot{x}(t) = G(x(t), \mathbf{S}(t))$ with the initial condition $x(0) = x_0$. Using the vertices of the template shape x_0 as initial conditions, the integration of this equation computes the deformation of the template shape $\mathbf{X}_0 = \mathbf{X}(0)$ to $\mathbf{X}(1)$ that is equal to $\phi(\mathbf{X}_0)$ with ϕ is the diffeomorphism parameterized by \mathbf{S}_0 .

2.3 Atlas Estimation

The N_{su} template-to-subject deformations in the criterion (1) are parameterized by N_{su} vectors \mathbf{S}_0^i . Each of these vectors serves as the initial condition in (2). Then, the template deformation is obtained by integration of (2) followed by (3). We choose to use the same set of control points \mathbf{c}_0 for all subjects, which defines a common basis for the deformations parameterization. By contrast, the set of initial momenta α_0^i are subject-specific. As a regularizer of the deformations, we use the energy of the set of particles $\sum_{p,q} \alpha_{0,p}^i {}^t K(c_{0,p}, c_{0,q}) \alpha_{0,q}^i$, which is also the geodesic distance connecting the template to the i^{th} shape [8]. In order to select the most relevant subset of control points, we add a L^1 penalty to (1), so that the criterion to minimize writes:

$$E(\mathbf{X}_0, \mathbf{c}_0, \{\alpha_0^i\}) = \sum_{i=1}^{N_{\text{su}}} \left\{ \frac{1}{2\sigma^2} D(\mathbf{X}_0^i(1), \mathbf{X}_i) + \sum_{p,q} \alpha_{0,p}^i {}^t K(c_{0,p}, c_{0,q}) \alpha_{0,q}^i + \gamma \sum_p \|\alpha_{0,p}^i\| \right\} \quad (4)$$

subject to:

$$\begin{cases} \dot{\mathbf{S}}^i(t) = F(\mathbf{S}^i(t)) & \text{with } \mathbf{S}^i(0) = \{\mathbf{c}_0, \alpha_0^i\} \\ \dot{\mathbf{X}}_0^i(t) = G(\mathbf{X}_0^i(t), \mathbf{S}^i(t)) & \text{with } \mathbf{X}_0^i(0) = \mathbf{X}_0 \end{cases} \quad (5)$$

The first equation in (5) is the equations of motion of the particles, like in (2). The second equation is deformation of the template parameterized by the particles motion, like in (3). σ^2 and γ balance the data term against the regularization terms. The variables to be optimized are: (i) the position of the vertices of the template shape \mathbf{X}_0 , (ii) the position of the control points in the template domain \mathbf{c}_0 and (iii) the N_{su} initial momenta α_0^i that parameterize each template-to-subject deformation. In practice, we also regularize the template shape defined by \mathbf{X}_0 by applying a penalty on Gaussian curvature of the mesh.

Only the first two terms in (4) are differentiable. As shown in the supplementary material accessible at the first author's webpage, the gradient of data term is given as:

$$\nabla_{\alpha_0^i} D = \xi^{\alpha,i}(0) \quad \nabla_{\mathbf{c}} D = \sum_{i=1}^{N_{\text{su}}} \xi^{c,i}(0) \quad \nabla_{\mathbf{X}_0} D = \sum_{i=1}^{N_{\text{su}}} \theta^i(0)$$

where the auxiliary variables $\xi^i(t) = \{\xi^{c,i}(t), \xi^{\alpha,i}(t)\}$ (of the same size as $\mathbf{S}^i(t)$) and $\theta^i(t)$ (of the same size as \mathbf{X}_0) satisfy the linear ODEs:

$$\begin{aligned} \dot{\theta}^i(t) &= -(\partial_1 G(\mathbf{X}_0^i(t), \mathbf{S}^i(t)))^t \theta^i(t) \text{ with } \theta^i(1) = \frac{1}{2\sigma^2} \nabla_{\mathbf{X}_0^i(1)} D(\mathbf{X}_0^i(1), \mathbf{X}_i) \\ \dot{\xi}^i(t) &= -\left(\partial_2 G(\mathbf{X}_0^i(t), \mathbf{S}^i(t))^t \theta^i(t) + d_{\mathbf{S}^i(t)} F(\mathbf{S}^i(t))^t \xi^i(t) \right) \text{ with } \xi^i(1) = 0 \end{aligned} \quad (6)$$

To compute the gradient, one integrates the flow equations (5) forward in time to build the deformations of the template shape. Then, one computes the

gradient of the residual data term ∇D , which serves as initial conditions in (6). The ODEs (6) transport this information from each subject’s space back to the template space, where the final value of the auxiliary variables $\theta(t)$ (resp. $\xi(t)$) is used to update the template (resp. the control point positions and momenta).

To optimize (4), which combines differentiable terms denoted E_{L^2} with an L^1 penalty, we use an adapted gradient-descent scheme called Fast Iterative Shrinkage and Thresholding Algorithm [1]. The template and control points are not affected by the L^1 term and are updated using a gradient-descent step at each iteration. By contrast, the momenta $\alpha_{0,p}^i$ are updated according to:

$$\alpha_{0,p}^i \leftarrow S_{\tau\gamma} \left(\left\| \alpha_{0,k}^i - \tau \nabla_{\alpha_{0,k}^i} E_{L^2} \right\| \right) \frac{\alpha_{0,k}^i - \tau \nabla_{\alpha_{0,k}^i} E_{L^2}}{\left\| \alpha_{0,k}^i - \tau \nabla_{\alpha_{0,k}^i} E_{L^2} \right\|}, \quad (7)$$

where τ is the current step-size of the gradient descent and S the usual soft-thresholding function $S_\lambda(x) = \max(0, x - \lambda) + \min(0, x + \lambda)$. This function zeroes out the momenta that are too small in magnitude, thus ensuring sparsity in the parametrization of deformations.

The parameters of the algorithm are the trade-offs σ and γ , the standard deviation of the Gaussian kernels for the momenta σ_V and the currents metric σ_W [15, 4]. We initialize the template shape with an ellipsoid for each connected component of the shapes, the control points with a regular lattice of step σ_V , and the initial momenta are set to zero.

3 Results

We apply our method to a study that seeks to compare neuroimaging, genetics, and neurotransmitter properties of patients with Down’s syndrome and healthy controls. We construct an atlas from surfaces of three different deep brain structures: amygdala, hippocampus, and putamen (Fig. 1). We initialize the atlas with one ellipsoid for each of the three anatomical structures, and we initialize the control points with a regular lattice of 650 points. After the optimization, the template shapes capture the common anatomical features across the populations (Fig. 2-left), and are given as meshes with the same topology as the initial set of ellipsoids. The parameterization of the template-to-subject deformations were also optimized: control points are moved toward the surfaces (where their theoretical optimal locations are, as shown in [15]) and the sparsity prior selects a subset of control points (99 out of 650) that carry a non-zero momentum vector. Results were generated using $\sigma = 10^{-4}$, $\gamma = 3 \times 10^6$, $\sigma_V = 5$, and $\sigma_W = 2$.

By comparison, we show the atlas built with the method of [4] in Fig. 3. The template shape is given as a set of disconnected normals: the continuity of the surfaces and the number of connected components of the input shapes have not been preserved. Moreover, the template-to-subjects deformations are parameterized with momenta located at the same place as the template normals. By contrast, our method optimized the position and the number of the control




Fig. 1. Sample input shapes where hippocampus, amygdala, and putamen are shown in yellow, blue, and cyan respectively.




Fig. 2. Atlas construction: template shape (left) and parameterization of the template-to-subject deformations (right). The template is initialized with an ellipsoid per connected components (top-left). After the optimization, the template shapes still have the same topology as the sample shapes. Simultaneously, template-to-subjects deformations are estimated, which are parameterized by the momentum vectors (red arrows). The vectors are located at the position of the control points, which are the same for all subjects. Control point positions were initialized as the nodes of a regular lattice, and our algorithm moves the control points to their optimal position near the surfaces and selects the most relevant ones according to the sparsity prior.

points independently of the vertices of the template shapes. Consequently, deformations are parameterized by more than 40 times fewer momenta. Constraining the template to remain a mesh has not introduced bias in the estimation: the norm of the difference between the templates generated by the two methods is 3.4×10^{-5} , which is much smaller than the standard deviation and below the usual threshold of 3 times standard deviation to decide statistical significance.

We construct a common template for the combined DS and healthy populations, and perform a Principal Component Analysis (PCA) on the momentum vectors for each population separately. The results in Fig. 4 show that the two populations contain different variability at different objects and at different locations within each object. The sparse momentum vectors enable statistical analysis in lower dimensional space which has great potential for clinical studies.




Fig. 3. Comparison of hippocampus template generated using [4] and our method. Arrows indicate the momenta driving the registration of the template to the first subject. Our method generates topologically correct mesh as opposed to triangle normals, and with momenta that are not constructed to be located at the surface triangles.




Fig. 4. Statistical analysis of momentum vectors for Down's syndrome (DS) and healthy control population. Left: template deformed according to the average momenta for DS (red) and control (yellow). Center and Right: Templates deformed according to the first mode of PCA on momentum vectors for DS and control population, demonstrating the different modes of variability. Colors indicate the magnitude of displacement from the population mean shape (in mm).

4 Conclusions

We propose a new method for estimating shape atlases using the currents metric. In contrast to [6, 4], the template shape is given as a mesh that has the same topology as the shapes in the population and the parameterization of the template-to-subject deformations were optimized independently of the template shape. We use a single gradient descent for estimating the template and the deformations, in contrast to the alternating minimization in [6, 4]. This scheme is much more efficient than using the matching pursuit technique and shows good convergence properties, even with the most naive initialization. Our non-parametric method makes use of the metric on currents, which enables the use of shapes without point correspondences, and therefore with minimal pre-processing.

The method also provides automatic parametrization of diffeomorphisms for mapping subjects within a population, where the parameters are constrained to lower dimensional spaces. We demonstrated the potential for these sparse parametrizations for performing statistical analysis on shapes from Down's syndrome and healthy control population groups, which can enable future research correlating brain function, anatomy, and neurocircuitry.

In the future, we plan to analyze the robustness and statistical power of the parametrizations provided by our method, in particular in High Dimension Low Sample Size (HDLSS) settings that are typical in imaging studies of populations.

Acknowledgements

This work has been supported by NIH grants U54 EB005149 (NA-MIC), 1R01 HD067731, 5R01 EB007688, 2P41 RR0112553-12.

References

1. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences* 2(1), 183–202 (2009)
2. Chui, H., Rangarajan, A., Zhang, J., Leonard, C.: Unsupervised learning of an atlas from unlabeled point-sets. *IEEE PAMI* 26(2), 160–172 (2004)
3. Durrleman, S., Fillard, P., Pennec, X., Trouvé, A., Ayache, N.: Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents. *NeuroImage* 55(3), 1073 – 1090 (2011)
4. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Statistical models of sets of curves and surfaces based on currents. *Med Image Anal* 13(5), 793–808 (2009)
5. Durrleman, S., Prastawa, M., Gerig, G., Joshi, S.: Optimal data-driven sparse parameterization of diffeomorphisms for population analysis. In: Székely, G., Hahn, H. (eds.) *IPMI 2011. LNCS*, vol. 6801, pp. 123–134 (2011)
6. Glaunès, J., Joshi, S.: Template estimation from unlabeled point set data and surfaces for computational anatomy. In: *Proc. of the International Workshop on the Mathematical Foundations of Computational Anatomy (MFCA-2006)* (2006)
7. Hufnagel, H., Pennec, X., Ehrhardt, J., Ayache, N., Handels, H.: Generation of a statistical shape model with probabilistic point correspondences and EM-ICP. *International Journal of CARS (IJCARS)* 2(5), 265–273 (2008)
8. Joshi, S., Miller, M.: Landmark matching via large deformation diffeomorphisms. *IEEE Trans. Image Processing* 9(8), 1357–1370 (2000)
9. Kurtek, S., Klassen, E., Ding, Z., Jacobson, S., Jacobson, J., Avison, M., Srivastava, A.: Parameterization-invariant shape comparisons of anatomical surfaces. *Trans. Med. Imag.* 30, 849–858 (2011)
10. Marsland, S., McLachlan, R.I.: A Hamiltonian particle method for diffeomorphic image registration. In: Karssemeijer, N., Lelieveldt, B.P.F. (eds.) *IPMI 2007. LNCS*, vol. 4548, pp. 396–407 (2007)
11. Pennec, X., Fillard, P., Ayache, N.: A Riemannian framework for tensor computing. *International Journal of Computer Vision* 66(1), 41–66 (2006)
12. Pizer, S., Fletcher, P., Joshi, S., Thall, A., Chen, J., Fridman, Y., Fritsch, D., Gash, A., Glotzer, J., Jiroutek, M., Lu, C., Muller, K., Tracton, G., Yushkevich, P., Chaney, E.: Deformable M-Reps for 3D Medical Image Segmentation. *International Journal of Computer Vision* 55(2-3), 85–106 (2003)
13. Székely, G., Kelemen, A., Brechbühler, C., Gerig, G.: Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible fourier contour and surface models. *Medical Image Analysis* 1, 19–34 (1996)
14. Vaillant, M., Miller, M., Younes, L., Trouvé, A.: Statistics on diffeomorphisms via tangent space representations. *NeuroImage* 23, 161–169 (2004)
15. Vaillant, M., Glaunès, J.: Surface matching via currents. In: Christensen, G.E., Sonka, M. (eds.) *IPMI 2005. LNCS*, vol. 3565, pp. 381–392 (2005)