Contents lists available at ScienceDirect

Medical Image Analysis

Editorial

Statistical shape analysis: From landmarks to diffeomorphisms

Miaomiao Zhang, Polina Golland*

Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, United States

ARTICLE INFO

ABSTRACT

We offer a blazingly brief review of evolution of shape analysis methods in medical imaging. As the representations and the statistical models grew more sophisticated, the problem of shape analysis has been gradually redefined to accept images rather than binary segmentations as a starting point. This transformation enabled shape analysis to take its rightful place in the arsenal of tools for extracting and understanding patterns in large clinical image sets. We speculate on the future developments in shape analysis and potential applications that would bring this mathematically rich area to bear on clinical practice.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

The study of anatomical shape changes and of their relationship with disease processes is one of the central topics in medical image analysis. For example, many neurodegenerative disorders cause widespread brain shape changes that can be observed in 3D brain MRI scans, as illustrated in Fig. 1. Identifying statistical shape differences between healthy subjects and patients affected by a disease promises to provide new clinical insights and ultimately improve diagnosis and treatment. In this paper, we briefly review the evolution of representations and statistical modeling techniques used for anatomical shape analysis and highlight recent developments that take us closer to routine applications of shape analysis based on large collections of medical images.

2. Shape representation

Shape representations commonly used in medical image analysis include landmarks, implicit representations, parametric representations, medial models, and deformation-based descriptors. Many variants have been proposed in the field based on these fundamental descriptors.

Landmarks are points on the object's boundary that can be identified reliably (Cootes et al., 1995; Bookstein, 1997). Landmarks are detected automatically or placed manually based on geometric properties of the surface, such as curvature, and image intensity. Since the quality of landmarks directly affects the statistical efficiency of the resulting shape model, substantial research effort has been focusing on selecting distinct and reproducible landmarks, including automatic methods for achieving compact representations of the group variability (Davies et al., 2001). Landmark-based methods most often characterize the object boundary and ignore the structure of the object's interior.

To overcome the sensitivity of the descriptor to the placement of the landmarks, a distance transform is commonly used as an implicit representation of the object's boundary (Leventon et al., 2000). While more robust to the noise in the boundary, such **implicit descriptors** give rise to a highly non-linear space, necessitating approximations when modeling shape distributions (Golland et al., 2005).

Parametric descriptors represent the object boundary in a particular functional basis and thus reduce the representation to a small number of discrete coefficients that capture the entire continuous boundary. Examples include the Fourier basis in 2D (Staib and Duncan, 1992) and the spherical harmonics in 3D (Brechbühler et al., 1995). While offering computational efficiency and robustness to noise in landmark locations, this inherently global descriptor makes it challenging to localize shape changes.

Based on the original work of Blum (1973), the **medial axis** descriptor represents the shape as an envelope of spheres whose centers define the object's "core" and whose radii determine the object's thickness at each point of the core. The original medial axis representation is sensitive to the noise in the binary label map that defines the object, which is overcome by robust fitting methods applied to medial representations of constrained topology (Pizer et al., 1999).

In contrast to the earlier descriptors that focused on binary segmentations of a specific organ, **deformation-based representations** capture shape information in the entire image (Christensen et al., 1993; Rueckert et al., 2003). With the underlying

* Corresponding author. Fax: +1 617-258-6287.
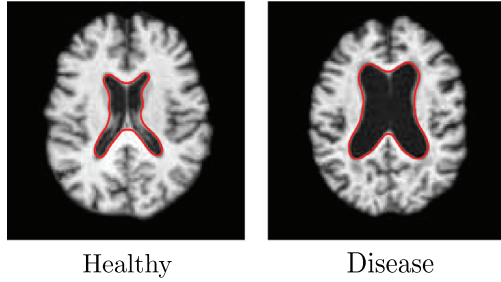*E-mail address:* polina@csail.mit.edu (P. Golland).

Fig. 1. Axial slices of example scans of a healthy subject and a patient from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database. The red contour highlights changes in the shape of the ventricles due to brain tissue atrophy in Alzheimer's disease.

assumption that the geometric information in the deformations conveys a shape, descriptors in this class arise naturally by matching a template to an input image with smoothness constraints on the deformation field. In many applications, it is natural to require the deformation to be a diffeomorphism, i.e., a differentiable, bijective mapping with a differentiable inverse. While biological motivation for using diffeomorphic deformations may exist in some applications, the biggest advantage of using diffeomorphisms is computational. Employing diffeomorphic transformations greatly facilitates the mapping between the atlas and subject coordinate systems. Moreover, the theoretical framework of Large Deformation Diffeomorphic Metric Mapping (LDDMM) defines a metric in the space of deformations that in turn induces a distance metric on the shape space (Beg et al., 2005).

The deformation-based approach to shape representation transformed shape analysis from requiring binary segmentations to operating on raw images. This change enabled many new applications of shape analysis for which no robust segmentation methods are available.

3. Statistical analysis

Principal component analysis (PCA) is widely used in combination with various shape representations to capture anatomical variability in a population (Cootes et al., 1995; Staib and Duncan, 1992; Leventon et al., 2000; Brechbühler et al., 1995). Many descriptors give rise to non-linear spaces, which require statistical analysis to account for the non-linearities in the model. Examples of such non-linear modeling methods include kernel PCA (Twining and Taylor, 2001), principal geodesic analysis (PGA) that estimates lower-dimensional geodesic subspaces by minimizing the sum-of-squared geodesic distances (Fletcher et al., 2003), and clustering (Sabuncu et al., 2009). And while most sophisticated methods for modeling high-dimensional distributions that capture low-dimensional phenomena has been attempted in the context of shape analysis, PCA and its manifold variants remain a workhorse of this field.

4. Deformation-based representations meet statistics on manifolds

In this section, we illustrate the ideas reviewed above in the context of a contemporary formulation based on a Bayesian variant of PGA. This particular approach encodes shape in the space of deformations parameterized through the so called geodesic shooting construction (Miller et al., 2006). The model employs an automatic relevance determination prior to encourage a compact representation. Detailed derivations of the statistical model and of the resulting inference algorithm can be found in (Zhang and Fletcher, 2015).
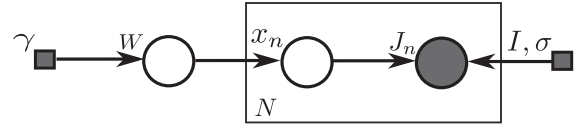

Fig. 2. Graphical representation of BPGA.

4.1. Diffeomorphisms for shape representation

We define diffeomorphic transformations on a d-dimensional torus $\Omega = \mathbb{R}^d / \mathbb{Z}^d$. The tangent space of diffeomorphisms at identity is the space V of smooth vector fields on Ω. For any two vector fields $v, w \in V$, there exists a weak Sobolev metric such that

$$\langle v, w \rangle_V = \int_\Omega (Lv(x), w(x)) \, dx, \tag{1}$$

where L is a symmetric, positive-definite, differential operator that induces a metric on the shape space. Diffeomorphisms are generated by flows of time-varying velocity fields as a solution to an ordinary differential equation

$$\frac{d\phi(t, x)}{dt} = v(t, \phi(t, x)), \tag{2}$$

where v and ϕ are the time-varying velocity field and the resulting diffeomorphic transformation, respectively.

Given an initial velocity field at time $t = 0$, the geodesic path on the manifold of diffeomorphisms is uniquely determined by the Euler-Poincaré equations (EPDiff) (Arnol'd, 1966; Miller et al., 2006):

$$\frac{dv}{dt} = -L^{-1} \big[(Dv)^T (Lv) + D(Lv) \, v + (Lv) \, \text{div} \, v \big], \tag{3}$$

where D is the Jacobian and div is the divergence operator. This process is known as geodesic shooting (Miller et al., 2006). This construction provides a representation of diffeomorphisms via initial velocities.

With a slight abuse of notation, in the remainder of the paper we use v to denote the initial velocity field and ϕ to denote the corresponding final deformation.

4.2. Image likelihood

Fig. 2 illustrates the graphical model for Bayesian principal geodesic analysis (BPGA), first presented in Zhang and Fletcher (2015). Given a set of N deformation fields $\{\phi_n\}$, we treat the input images $\{J_n\}$ as iid samples from a Gaussian distribution induced by the model of the image noise ($n = 1, \cdots, N$). In particular, input image J_n is generated by applying deformation ϕ_n to the image template I:

$$p(J_n \mid \phi_n; I, \sigma^2) \propto \exp \left(-\frac{1}{2\sigma^2} \left\| J_n - I \circ \phi_n^{-1} \right\|_{L^2}^2 \right), \tag{4}$$

where σ^2 is the image noise variance. The norm inside the exponent is the $L^2(\Omega, \mathbb{R})$ norm defined on a finite discretized grid. More sophisticated models of image noise can be readily converted into a likelihood model for the observed images.

4.3. Generative model of shape

We let the columns of matrix $W \in \mathbb{R}^{d \times K}$ represent the K principal modes of variation in the d-dimensional space of the initial velocity fields. We let $x_n \in \mathbb{R}^K$ denote a vector of K latent loading weights such that the initial velocity field $v_n = W x_n$ generates diffeomorphism ϕ_n in Eq. (4). Similar to the Bayesian variant of
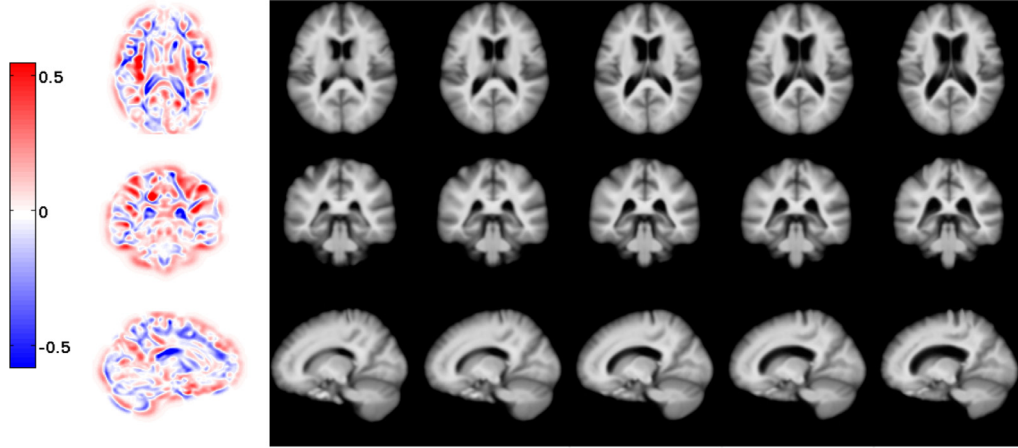
Fig. 3. Top to bottom: axial, coronal and sagittal views of shooting the estimated image template I along the first principal mode W_1. Left to right: log determinant of the Jacobian at $-3\gamma_1^{-1}W_1$, resulting image at $\{-3, -1.5, 0, 1.5, 3\} \times \gamma_1^{-1}W_1$ respectively. (For interpretation of the references to color in the text, the reader is referred to the web version of this article.)

PCA (Bishop, 1999), we impose a Gaussian prior on the loading factors:

$$p(x_n \,|\, W) \propto \exp\left(-\frac{1}{2}\|Wx_n\|_V^2\right). \tag{5}$$

In contrast to Bayesian PCA, this distribution includes principal components W to ensure the smoothness of the resulting geodesic shooting path. The prior distribution on W encourages sparsity by suppressing small principal initial velocity values towards zero, thus providing dimensionality reduction through automatic relevance determination (Bishop, 1999):

$$p(W; \gamma) \propto \prod_{k=1}^{K} \left(\frac{\gamma_k}{2\pi}\right)^{\frac{d}{2}} \exp\left(-\frac{\gamma_k}{2}\|W_k\|_V^2\right), \tag{6}$$

where W_k is the kth principal component of the model and γ_k is the precision of the principal component W_k.

4.4. Learning and inference

The maximum likelihood estimates of the image template I, the noise variance σ^2, and the precision hyperparameters $\{\gamma_k\}$ are obtained through an iterative algorithm that is derived by observing that maximizing the likelihood of the input images is equivalent to minimizing the cost function that leads to geodesic shooting (Miller et al., 2006). The optimization procedure produces estimates of the hyperparameters $\{\gamma_k\}$ that often tend to infinity. Since high precision is equivalent to concentration around zero, this process yields sparse solutions. Other choices of sparsity, such as Laplacian priors, can also be used to implement dimensionality reduction. Once the algorithm converges, the maximum *a posteriori* estimates of the latent loading weights x_n provide a compact representation of the anatomical shape in image J_n and can be used for further analysis in clinical studies.

Fig. 3 illustrates the application of this approach in a study of brain shape variability based on a set of 3D MRI scans of healthy subjects in ADNI. We display regions of expansion (in red) and contraction (in blue) captured in the first principal component, as well as the anatomical variability represented by this component. We observe that ventricle size change is a dominant factor in brain shape variability in this cohort.

5. Conclusions and discussion

In this paper, we briefly reviewed the evolution of statistical shape representations, from landmarks to diffeomorphisms, and of the statistical methods that have been employed to capture the variability in the resulting shape descriptors. We illustrated a contemporary approach to shape analysis by demonstrating one possible combination of diffeomorphism-based shape representation coupled with the Bayesian principal geodesic analysis to extract principal modes of variation in the brain anatomy of healthy aging population.

In addition to ever increasing sophistication in the mathematical methods for capturing shape and its variation, the field underwent an exciting development of realizing shape analysis on raw images rather than requiring accurate segmentation as a prerequisite step. While this transformation alleviated the challenges of obtaining high quality segmentations, it came at a cost of computational complexity. Since deformation-based shape descriptors exist in a high dimensional space defined by the image grid, manipulating these descriptors requires substantial computational resources and specialized programming techniques. We expect the future turns of the spiral in shape analysis evolution to produce efficient deformation representations that still capture all the intricate and clinically relevant anatomical detail. As an additional benefit, such representations will also bring us closer to efficient implementations of image alignment and practical applications of shape analysis in large scale studies of anatomical change in development, aging and disease.

We conclude by observing that shape analysis has not yet made its way into the standard medical practice or clinical research. While technical challenges of characterizing complex anatomical shape have been successfully addressed by the evolution of methods in medical image computing, the challenge of translation remains for the current and future researchers in this area. Only when clinical researchers and practitioners rely on shape analysis tools in the everyday decision making, can we confidently state that the research reviewed in this paper has lived up to its true potential.

Acknowledgments

References

Arnol'd, V.I., 1966. Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. Ann. Inst. Fourier 16, 319–361.

Beg, M., Miller, M., Trouvé, A., Younes, L., 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. Int. J. Comput. Vis. 61 (2), 139–157.

Bishop, C.M., 1999. Bayesian PCA. Advances in Neural Information Processing Systems 382–388.

Blum, H., 1973. Biological shape and visual science (part I). J. Theor. Biol. 38 (2), 205–287.

Bookstein, F.L., 1997. Morphometric Tools for Landmark Data: Geometry and Biology. Cambridge University Press.

Brechbühler, C., G., G., Kübler, O., 1995. Parametrization of closed surfaces for 3-d shape description. Comput. Vis. Image Understand. 61 (2), 154–170.

Christensen, G.E., Rabbitt, R.D., Miller, M.I., 1993. A deformable neuroanatomy textbook based on viscous fluid mechanics. In: 27th Annual Conference on Information Sciences and Systems. Citeseer, pp. 211–216.

Cootes, T.F., Taylor, C.J., Cooper, D.H., Gr aham, J., et al., 1995. Active shape models-their training and application. Comput. Vis. Image understand. 61 (1), 38–59.

Davies, R.H., Twining, C.J., Cootes, T.F., Waterton, J.C., Taylor, C.J., 2001. A minimum description length approach to statistical shape modelling. IEEE Trans. Med. Imag. 21, 525–537.

Fletcher, P.T., Lu, C., Joshi, S., 2003. Statistics of shape via principal geodesic analysis on Lie groups. In: Computer Vision and Pattern Recognition, vol. 1. IEEE, pp. I–95.

Golland, P., Grimson, W.E.L., Shenton, M.E., Kikinis, R., 2005. Detection and analysis of statistical differences in anatomical shape. Med. Image Anal. 9 (1), 69–86.

Leventon, M.E., Grimson, W.E.L., Faugeras, O., 2000. Statistical shape influence in geodesic active contours. In: Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on, vol. 1. IEEE, pp. 316–323.

Miller, M.I., Trouvé, A., Younes, L., 2006. Geodesic shooting for computational anatomy. J. Math. Imag. Vis. 24 (2), 209–228.

Pizer, S.M., Fritsch, D.S., Yushkevich, P.A., Johnson, V.E., Chaney, E.L., 1999. Segmentation, registration, and measurement of shape variation via image object shape. Med. Imag. IEEE Trans. 18 (10), 851–865.

Rueckert, D., Frangi, A.F., Schnabel, J.A., 2003. Automatic construction of 3-d statistical deformation models of the brain using nonrigid registration. Med. Imag. IEEE Trans. 22 (8), 1014–1025.

Sabuncu, M.R., Balci, S.K., Shenton, M.E., Golland, P., 2009. Image-driven population analysis through mixture modeling. Med. Imag. IEEE Trans. 28 (9), 1473–1487.

Staib, L., Duncan, J., 1992. Boundary finding with parametrically deformable models. IEEE Trans. Pattern Anal. Mach. Intell. 14 (11), 1061–1075.

Twining, C.J., Taylor, C.J., 2001. Kernel principal component analysis and the construction of non-linear active shape models. In: British Machine Vision Conference, pp. 1–10.

Zhang, M., Fletcher, P.T., 2015. Bayesian principal geodesic analysis for estimating intrinsic diffeomorphic image variability. Med. Image Anal. 25 (1), 37–44.