Fast Geodesic Regression for Population-Based Image Analysis

Yi Hong[1(✉)], Polina Golland[2], and Miaomiao Zhang[2]

[1] Computer Science Department, University of Georgia, Athens, USA
yi.hong@uga.edu
[2] Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, USA

Abstract. Geodesic regression on images enables studies of brain development and degeneration, disease progression, and tumor growth. The high-dimensional nature of image data presents significant computational challenges for the current regression approaches and prohibits large scale studies. In this paper, we present a fast geodesic regression method that dramatically decreases the computational cost of the inference procedure while maintaining prediction accuracy. We employ an efficient low dimensional representation of diffeomorphic transformations derived from the image data and characterize the regressed trajectory in the space of diffeomorphisms by its initial conditions, i.e., an initial image template and an initial velocity field computed as a weighted average of pairwise diffeomorphic image registration results. This construction is achieved by using a first-order approximation of pairwise distances between images. We demonstrate the efficiency of our model on a set of 3D brain MRI scans from the OASIS dataset and show that it is dramatically faster than the state-of-the-art regression methods while producing equally good regression results on the large subject cohort.

1 Introduction

In medical research, image time-series are collected for individual subjects or in a population to monitor and study aging, disease progression, brain development and degeneration. For instance, brain magnetic resonance imaging (MRI) scans capture anatomical and functional changes in individual brains. Summarizing the characteristic patterns of these changes will improve our understanding of brain functions and disease progression for developing early diagnosis and effective treatment.

Image regression has been commonly used to estimate such changes. Existing approaches include piecewise regression based on image registration [5,10], methods based on kernel regression [2,4], geodesic regression [11,12], polynomial regression [6], and spline regression [13]. In contrast to piecewise regression and kernel regression approaches, geodesic and higher-order regression methods aim to estimate a parametric model that minimizes the sum of squared distances between the observed images and the corresponding images on the regression trajectory. The compact representation in the form of model parameters is then

M. Descoteaux et al. (Eds.): MICCAI 2017, Part I, LNCS 10433, pp. 317–325, 2017.
DOI: 10.1007/978-3-319-66182-7_37

used in further statistical analysis, e.g., performing group comparisons between patients with a particular disease and normal controls and identifying statistical differences between such cohorts. Unfortunately, closed-form solutions generally do not exist for this problem, especially in the case of diffeomorphisms [11]. The distance metric defined by diffeomorphic image registration [3] is too expensive to compute for an iterative optimization in the high dimensional image space. Previous work on first order approximations have been used to derive a closed-form solution for the distance metric [8]. However, all the computations were still implemented on a full dense image grid, which typically requires massive amounts of time and memory.

In this paper, we derive a fast geodesic regression method that utilizes a finite dimensional Fourier representation of the tangent space of diffeomorphisms [15] to enable efficient estimation of geodesic regression for image time-series. We define a distance metric in the low dimensional bandlimited space that leads to a fast estimation of the regression trajectory, which includes the initial image ('intercept') and the initial velocity field ('slope'). Moreover, we apply the first-order approximation [8] to the distance metric and derive a closed-form solution that eliminates the need for iterative estimation. In particular, our model reduces the optimization problem of image regression to a collection of independent pairwise image registrations that is easily implemented in parallel. The resulting initial velocity is computed as a weighted average of the velocity fields in the low dimensional Fourier space estimated from pairwise registrations. We demonstrate the efficiency of our model on a set of 3D brain MRIs from the OASIS dataset and show that it is dramatically faster than the existing regression method [8] without sacrificing accuracy.

2 Background: Geodesic Regression

We first review the geodesic regression for image time-series in the setting of large deformation diffeomorphic metric mapping (LDDMM) with geodesic shooting [11]. Assume at P time instants $\{t_i\}(i = 1, \cdots, P)$, we have a set of images $\{Y_{ij}\}(j = 1, \cdots, N_i)$, that is, at each time instant t_i there are N_i images. The problem of geodesic regression is then formulated by minimizing an energy function

$$E(I_0, v_0) = \frac{1}{2}(\mathcal{L}v_0, v_0) + \frac{1}{\sigma^2}\sum_{i=1}^{P}\sum_{j=1}^{N_i}\text{Dist}(I_0 \circ \phi_i^{-1}, Y_{ij}), \tag{1}$$

where I_0 and v_0 are the unknown initial image ('intercept') and the unknown initial velocity field ('slope') that parameterize the regression geodesic at a starting time point t_0, σ^2 is a constant representing the noise variance, Y_{ij} is the jth image at time point t_i, and $d\phi_i/dt = (\Delta t_i v_0) \circ \phi_i$ with $\Delta t_i = t_i - t_0$. Here, (m, v) is a pairing of a velocity field v and its dual $m = \mathcal{L}v$, with a symmetric positive-definite differential operator $\mathcal{L}$. The distance function $\text{Dist}(\cdot, \cdot)$ measures the squared difference between the image Y_{ij} and its corresponding image

on the regression geodesic, i.e., the image I_0 deformed by the transformation ϕ_i. A commonly-used LDDMM framework [3] defines the distance metric

$$\text{Dist}(\mathcal{I}_1, \mathcal{I}_2) = \frac{1}{2}\int_0^1 (\mathcal{L}u_\tau, u_\tau)\, d\tau + \lambda \|\mathcal{I}_1 \circ \psi_1^{-1} - \mathcal{I}_2\|_2^2, \tag{2}$$

where $\mathcal{I}_1$, $\mathcal{I}_2$ are the images of interest, u_τ ($\tau \in [0,1]$) is the time-varying velocity field. The path of deformation fields ψ_τ is generated by $d\psi_\tau/d\tau = u_\tau \circ \psi_\tau$. Deformation $\psi_0 = \text{Id}$ is the identity element and λ is a positive weight parameter.

The geodesic shooting algorithm estimates the initial velocity field u_0 at $\tau = 0$ and relies on the fact that a geodesic path of transformations ψ_τ with a given initial condition u_0 can be uniquely determined through integrating the Euler-Poincaré differential equation (EPDiff) [1,9] as

$$\frac{\partial u_\tau}{\partial \tau} = -\mathcal{K}\left[(Du_\tau)^T m_\tau + Dm_\tau\, u_\tau + m_\tau \,\text{div}(u_\tau)\right], \tag{3}$$

where div is the divergence operator, D denotes Jacobian matrix, and $\mathcal{K} = \mathcal{L}^{-1}$ is the inverse of the smoothness operator $\mathcal{L}$ in (2). Since the optimal transformation ψ_1 can be parameterized by the given initial velocity u_0, we use an exponential map Exp to simply denote the relationship between u_0 and ψ_1 as $\psi_1 = \text{Exp}_{\text{Id}}(u_0)$.

The problem of minimizing the distance metric (2) can be equivalently reduced to optimizing over the initial velocity u_0 with the EPDiff (3). The distance metric can then be rewritten as

$$\text{Dist}(\mathcal{I}_1, \mathcal{I}_2) = \frac{1}{2}(\mathcal{L}u_0, u_0) + \lambda \|\mathcal{I}_1 \circ \psi_1^{-1} - \mathcal{I}_2\|_2^2, \quad \text{s.t. Eq. (3).} \tag{4}$$

In practice, we would have an inexact matching when measuring the distance due to the noise and appearance changes[1].

3 Fast Geodesic Regression

The standard iterative minimizing procedure for energy function (1) requires gradient computations in the high dimensional image space. To reduce its computational cost, a simple geodesic regression method [8] that derives a closed-form solution for the initial velocity v_0 by employing a first order approximation. All the computations are still implemented on the full dense image grid, which limits the model's applicability for large scale population analysis. In this paper, we introduce a fast geodesic regression algorithm that adopts the low dimensional bandlimited representation of the velocity fields in the Fourier space [15].

We define our model in the finite-dimensional Fourier space $\tilde{V}$ that represents bandlimited velocity fields with conjugate frequencies. Let $\mathcal{F}$ be a function that maps an element $u \in V$ from the image domain to the frequency domain $\tilde{u} \in \tilde{V}$, i.e., $\tilde{u} = \mathcal{F}[u]$. The inverse Fourier transform $\mathcal{F}^{-1}$ maps the signal $\tilde{u}$ back to the

[1] A metamorphosis approach [7,14] can be used instead to produce an exact matching.

image domain, i.e., $u = \mathcal{F}^{-1}[\tilde{u}]$. The correspondence between the initial velocity field $\tilde{u}_0$ and its associated transformation ψ_1 at time point $\tau = 1$ in the image domain is $\psi_1 = \mathrm{Exp}_{\mathrm{Id}}(\mathcal{F}^{-1}[\tilde{u}_0])$. With a slight abuse of notation, we drop the time index of the initial velocity $\tilde{u}_0$ and the deformation ψ_1 and use $\tilde{u}$ and ψ in the remainder of this paper.

Similar to [8], we define a first order approximation of the diffeomorphic distances in Eq. (1) between images through the exponential map in Fourier space. Suppose the initial velocity $\tilde{u}_{ij}$ and the corresponding diffeomorphic transformation ψ_{ij} map the initial image I_0 to the observed image Y_{ij}. Based on the rule of right composition, we have a transformation $\Phi_{ij} = \psi_{ij} \circ (\phi_i)^{-1}$ that measures the mapping error between the deformed initial image $\hat{I}_i = I_0 \circ \phi_i^{-1}$ and the observed image Y_{ij}. We now rewrite Φ_{ij} in the form of an exponential map:

$$\Phi_{ij} = \psi_{ij} \circ (\phi_i)^{-1} = \mathrm{Exp}_{\mathrm{Id}}(\mathcal{F}^{-1}[\tilde{u}_{ij}]) \circ \mathrm{Exp}_{\mathrm{Id}}(-\mathcal{F}^{-1}[\Delta t_i \tilde{v}_0]). \tag{5}$$

By taking the first-order approximation of Φ_{ij} in Eq. (5), we obtain

$$\Phi_{ij} \approx \mathrm{Exp}_{\mathrm{Id}}(\mathcal{F}^{-1}[\tilde{u}_{ij} - \Delta t_i \tilde{v}_0]),$$

which provides an approximation of the distance between images in terms of Fourier representations of the pairwise initial velocity fields $\tilde{u}_{ij}$ and the regression initial velocity $\tilde{v}_0$.

The minimal distance in Eq. (4) is empirically equivalent to the Riemannian length of the optimal transformation Φ_{ij}^* that corresponds to the optimal initial velocity $\tilde{v}_0^*$:

$$\begin{aligned} \mathrm{Dist}(\hat{I}_i, Y_{ij}) &= \|\Phi_{ij}^*\|^2 + \lambda \|\hat{I}_i \circ {\Phi_{ij}^*}^{-1} - Y_{ij}\|^2 \\ &\approx \|\mathrm{Exp}_{\mathrm{Id}}(\mathcal{F}^{-1}[\tilde{u}_{ij} - \Delta t_i \tilde{v}_0^*])\|^2, \end{aligned} \tag{6}$$

when $\hat{I}_i \circ {\Phi_{ij}^*}^{-1} \approx Y_{ij}$. Since the minimal-energy curve preserves constant speed along geodesics, we obtain

$$\begin{aligned} \|\mathrm{Exp}_{\mathrm{Id}}(\mathcal{F}^{-1}[\tilde{u}_{ij} - \Delta t_i \tilde{v}_0^*])\|^2 &= \frac{1}{2}\left(\mathcal{F}^{-1}[\tilde{\mathcal{L}}(\tilde{u}_{ij} - \Delta t_i \tilde{v}_0^*)], \mathcal{F}^{-1}[\tilde{u}_{ij} - \Delta t_i \tilde{v}_0^*]\right) \\ &= \frac{1}{2}\left(\tilde{\mathcal{L}}(\tilde{u}_{ij} - \Delta t_i \tilde{v}_0^*), \tilde{u}_{ij} - \Delta t_i \tilde{v}_0^*\right). \end{aligned} \tag{7}$$

Here, $\tilde{\mathcal{L}}$ is the Fourier representation of the differential operator $\mathcal{L}$, i.e., a d-dimensional Laplacian operator $(-\alpha\Delta + I)^c$ with a positive weight parameter α and a smoothness parameter c:

$$\tilde{\mathcal{L}}(\xi_1, \ldots, \xi_d) = (-2\alpha \sum_{q=1}^{d} (\cos(2\pi\xi_q) - 1) + 1)^c,$$

where $\xi_q (q = 1, \cdots, d)$ denotes frequency.

Using Eqs. (6) and (7), we are ready to approximate the geodesic regression formulation in Eq. (1) as

$$E(I_0, \tilde{v}_0) = \frac{1}{2}(\tilde{\mathcal{L}}\tilde{v}_0, \tilde{v}_0) + \frac{1}{2\sigma^2}\sum_{i=1}^{P}\sum_{j=1}^{N_i}\left(\tilde{\mathcal{L}}(\tilde{u}_{ij} - \Delta t_i\tilde{v}_0), \tilde{u}_{ij} - \Delta t_i\tilde{v}_0\right). \quad (8)$$

The initial image I_0 can be computed simultaneously with the registration-based initial velocities $\tilde{u}_{ij}$ by the unbiased atlas building algorithm [15].

Differentiating Eq. (8) w.r.t. the initial velocity $\tilde{v}_0$, we obtain

$$\nabla_{\tilde{v}_0} E = \tilde{v}_0 + \frac{1}{\sigma^2}\sum_{i=1}^{P}\sum_{j=1}^{N_i}\Delta t_i(\Delta t_i\tilde{v}_0 - \tilde{u}_{ij}) = 0.$$

Note that we compute the gradient in the Sobolev space by applying operator $\tilde{\mathcal{K}}$ for numerical stability, which cancels out the differential operator $\tilde{\mathcal{L}}$.

Finally, we arrive at a closed-form solution for $\tilde{v}_0$ as

$$\tilde{v}_0 = \Big(\sum_{ij}\Delta t_i\tilde{u}_{ij}\Big)/\Big(\sigma^2 + \sum_i N_i\Delta t_i^2\Big). \quad (9)$$

To summarize, our algorithm estimates the initial image I_0 and the velocity fields $\tilde{u}_{ij}$ jointly in a low dimensional bandlimited space by employing an efficient atlas building approach [15]. The initial velocity $\tilde{v}_0$ is then updated as a weighted average of the estimated velocity fields $\tilde{u}_{ij}$ as shown in Eq. (9). Since I_0 is the mean image of the entire population, the starting time point of the regression line in our paper is associated with the average of all time points as $\frac{1}{\sum_i^P N_i}\sum_{i=1}^{P} N_i t_i$.

4 Results

Data. We demonstrate our fast geodesic regression (FGR) model on a set of 3D brain MRIs of 129 subjects from the OASIS database, aged from 60 to 98. The MR scans were collected from 69 healthy controls (182 MR scans in total) and 60 individuals with mild to moderate Alzheimer's disease (AD) (136 MR scans in total). Each individual was scanned at 2-5 time points with the same resampled resolution $128 \times 128 \times 128$ and the voxel size of $1.25 \times 1.25 \times 1.25\,\mathrm{mm}^3$. All images underwent down-sampling, skull-stripping, intensity normalization to the range $[0, 1]$, bias field correction, and co-registration with affine transformations.

Experiments. We estimate the regression trajectory for each cohort separately. We use 16^3 dimensions for the Fourier representations of the velocity fields for patients and control group [15]. We set $\lambda = 1.0, \alpha = 3.0, c = 3.0$ for the operator $\tilde{\mathcal{L}}$ and the number of steps of the time integration in pairwise diffeomorphic image registration in the Fourier space is set to 10. We set $\sigma^2 = 0.01$ to balance the regularization and data matching error. We initialize I_0 as the average of

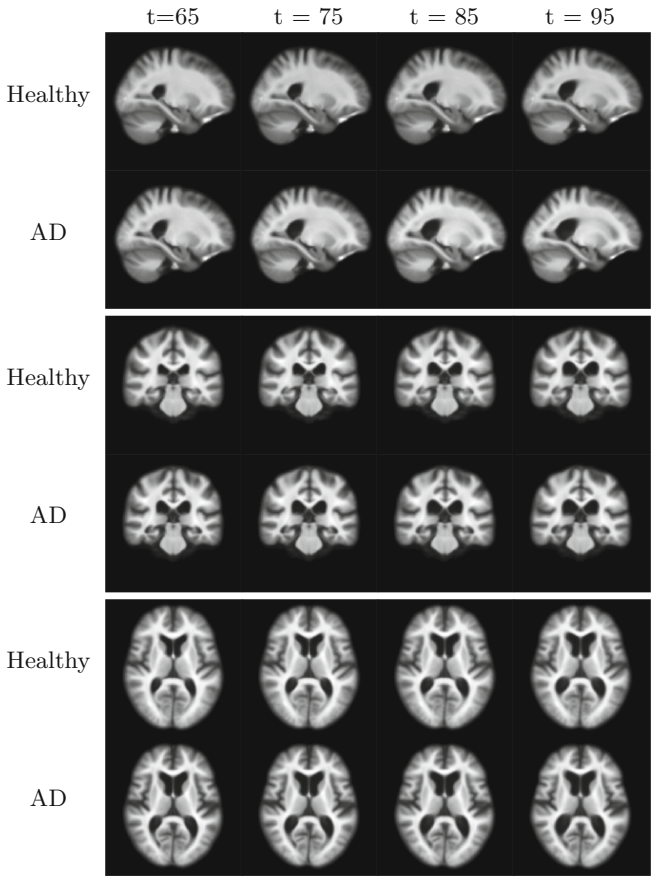

Fig. 1. Images at time points $t = 65, 75, 85, 95$ obtained via forward/backward shooting the estimated initial image by the estimated initial velocity. Sagittal (top), coronal (middle), and axial (bottom) views are shown for healthy subjects and AD patients.

image intensities and $\tilde{v}_0$ is a zero vector field. We evaluate accuracy, runtime, and memory consumption of our FGR model.

To evaluate the model's ability to capture the group trend, we compute the sum-of-squared-differences (SSD) between another 16 test image scans and the images obtained by shooting the image template I_0 with the estimated velocity field $\tilde{v}_0$ for the corresponding age of the subject. We employ the simple geodesic regression (SGR) [8] estimated on the full image grid as the baseline. For fair comparison, we keep all the parameters the same for both methods. We perform the two-sample hypothesis test on the computed SSDs between our method (FGR) and the baseline algorithm (SGR).

Experimental Results. Figure 1 visualizes the estimated group trends with shooting results for $t_i = 65, 75, 85, 95$ and illustrates the expansion and shrinkage

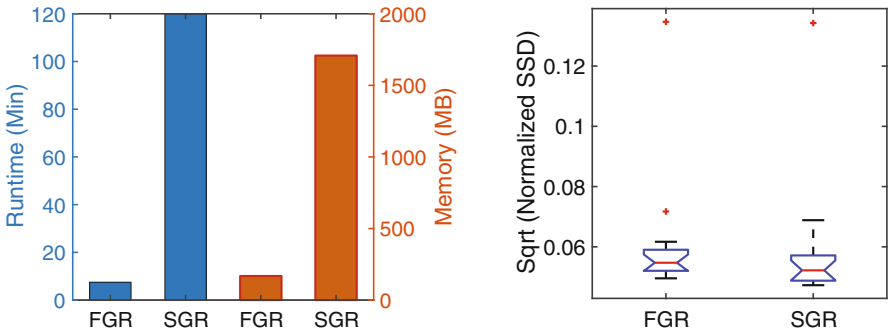

Fig. 2. Comparison of our fast geodesic regression (FGR) with the simple geodesic regression (SGR) [8]. Left: Runtime in minutes and memory in MB. Right: Normalized sum-of-squared-differences (SSD) between a test image and the corresponding image predicted by the model. There are 16 images in the test data set. The test image corresponding to the outlier on the top has quite difference image appearance from other images.

of the brain anatomy. The group trajectories for both cohorts in our study show anatomical changes consistent with [12], e.g., the changes in ventricle size is the dominant source of variability in both populations, while the dementia group has a faster degeneration rate.

As shown in the left panel of Fig. 2, the entire inference procedure finishes in 7.5 min with 168.4 MB memory by using parallel computing, while the simple geodesic regression method in a high dimensional image space [8] requires more than 2 h and 1708.1 MB memory using the same number of cores. The right panel of Fig. 2 reports the square root of the normalized SSD for our method (FGR) and the baseline algorithm (SGR). We therefore conclude that the two algorithms produce comparable results in terms of quality of image prediction of SSD errors. The difference in SSD is not statistically significant in a paired t-test ($p = 0.7391$).

5 Discussion and Conclusions

In this paper, we proposed a fast geodesic regression method that dramatically decreases the computational costs while offering comparable accuracy. We employ an efficient low dimensional representation of diffeomorphisms derived from the image data and estimate the regressed trajectory with a closed-form solution. Compared to the simple geodesic regression in [8], our approach is an order of magnitude faster and requires much less memory, with comparable accuracy for the estimated regression trajectories. This paves an efficient way to develop hierarchical regression model that processes large longitudinal datasets.

The group-level trend estimated by our method can be parallel transported to an individual's baseline image to further predict following-up scans; however, parallel transport is non-trivial in a high-dimensional image space and will

be considered in future work. Another promising direction is to generalize our model to polynomial regression that captures non-linear anatomical changes, for instance, the saturation effects in the aging brain.

Acknowledgments. This work was supported by NIH NIBIB NAC P41EB015902, NIH NINDS R01NS086905, and Wistron Corporation.

References

1. Arnol'd, V.I.: 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 (1966)
2. Banerjee, M., Chakraborty, R., Ofori, E., Okun, M.S., Viallancourt, D.E., Vemuri, B.C.: A nonlinear regression technique for manifold valued data with applications to medical image analysis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4424–4432 (2016)
3. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. Int. J. Comput. Vision **61**(2), 139–157 (2005)
4. Davis, B.C., Fletcher, P.T., Bullitt, E., Joshi, S.: Population shape regression from random design data. Int. J. Comput. Vision **90**(2), 255–266 (2010)
5. Durrleman, S., Pennec, X., Trouvé, A., Gerig, G., Ayache, N.: Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5761, pp. 297–304. Springer, Heidelberg (2009). doi:10.1007/978-3-642-04268-3_37
6. Hinkle, J., Fletcher, P.T., Joshi, S.: Intrinsic polynomials for regression on riemannian manifolds. J. Math. Imaging Vis. **50**(1–2), 32–52 (2014)
7. Hong, Y., Joshi, S., Sanchez, M., Styner, M., Niethammer, M.: Metamorphic geodesic regression. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012. LNCS, vol. 7512, pp. 197–205. Springer, Heidelberg (2012). doi:10.1007/978-3-642-33454-2_25
8. Hong, Y., Shi, Y., Styner, M., Sanchez, M., Niethammer, M.: Simple geodesic regression for image time-series. In: Dawant, B.M., Christensen, G.E., Fitzpatrick, J.M., Rueckert, D. (eds.) WBIR 2012. LNCS, vol. 7359, pp. 11–20. Springer, Heidelberg (2012). doi:10.1007/978-3-642-31340-0_2
9. Miller, M.I., Trouvé, A., Younes, L.: Geodesic shooting for computational anatomy. J. Math. Imaging Vis. **24**(2), 209–228 (2006)
10. Niethammer, M., Hart, G.L., Zach, C.: An optimal control approach for the registration of image time-series. In: Proceedings of the 48th IEEE Conference on Decision and Control, 2009 held Jointly with the 2009 28th Chinese Control Conference, CDC/CCC 2009, pp. 2427–2434. IEEE (2009)
11. Niethammer, M., Huang, Y., Vialard, F.-X.: Geodesic regression for image time-series. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011. LNCS, vol. 6892, pp. 655–662. Springer, Heidelberg (2011). doi:10.1007/978-3-642-23629-7_80
12. Singh, N., Hinkle, J., Joshi, S., Fletcher, P.T.: Hierarchical geodesic models in diffeomorphisms. Int. J. Comput. Vision **117**(1), 70–92 (2016)
13. Singh, N., Vialard, F.X., Niethammer, M.: Splines for diffeomorphisms. Med. Image Anal. **25**(1), 56–71 (2015)

14. Trouvé, A., Younes, L.: Metamorphoses through lie group action. Found. Comput. Math. **5**(2), 173–198 (2005)
15. Zhang, M., Fletcher, P.T.: Finite-dimensional lie algebras for fast diffeomorphic image registration. In: Ourselin, S., Alexander, D.C., Westin, C.-F., Cardoso, M.J. (eds.) IPMI 2015. LNCS, vol. 9123, pp. 249–260. Springer, Cham (2015). doi:10.1007/978-3-319-19992-4_19