Geodesic image regression with a sparse parameterization of diffeomorphisms

James Fishbaugh[1], Marcel Prastawa[1], Guido Gerig[1], and Stanley Durrleman[2]

[1] Scientific Computing and Imaging Institute, University of Utah, USA
[2] INRIA/ICM, Pitié Salpêtrière Hospital, Paris, France

Abstract. Image regression allows for time-discrete imaging data to be modeled continuously, and is a crucial tool for conducting statistical analysis on longitudinal images. Geodesic models are particularly well suited for statistical analysis, as image evolution is fully characterized by a baseline image and initial momenta. However, existing geodesic image regression models are parameterized by a large number of initial momenta, equal to the number of image voxels. In this paper, we present a sparse geodesic image regression framework which greatly reduces the number of model parameters. We combine a control point formulation of deformations with a L^1 penalty to select the most relevant subset of momenta. This way, the number of model parameters reflects the complexity of anatomical changes in time rather than the sampling of the image. We apply our method to both synthetic and real data and show that we can decrease the number of model parameters (from the number of voxels down to hundreds) with only minimal decrease in model accuracy. The reduction in model parameters has the potential to improve the power of ensuing statistical analysis, which faces the challenging problem of high dimensionality.

1 Introduction

A prominent tool for longitudinal analysis in medicine is regression of imaging data. Regression allows data from different subjects to be brought into temporal alignment, including images acquired at different time points and with different number of observations. Furthermore, regression is a necessary tool to align imaging data with cognitive assessments not necessarily acquired at the same time. Many longitudinal analysis frameworks require to compare such individual trajectories across populations [4,6,8,9,12]. Several methods for regression have been proposed, including the extension of piecewise linear regression to shapes represented as currents [4]. In [5,14] second-order models have been proposed which ensure that trajectories of shape change are smooth in time. Kernel regression was extended to image time series in [2]. These models require to store model parameters over time in order to characterize the evolution. This makes their use for statistical analysis cumbersome, as statistics are most naturally performed when model parameters can be compared in a common space.

This problem is addressed in the work of [7,11,13] where geodesic regression is proposed for image time series as well as in the Riemannian setting. The geodesic

regression model is generative, being fully characterized by the baseline image (the intercept) and the momenta vectors defining the geodesic at the baseline image (the slope). While these models cannot match the observed data as closely as the class of non-generative models discussed above (in general), the power of such models lies in the simplification of the statistical analysis. To compare two trajectories of change, one only needs to compare the initial conditions [10].

Although model parameters are located only at a baseline time point, the geodesic regression models of [11,13] still require to store a large number of parameters. In fact, the number of parameters for these geodesic regression models is equal to the number of image voxels, which can quickly become unwieldy, particularly for 3D medical images. However, the dynamics of image evolution are likely characterized by considerably fewer parameters. Intuitively, momenta should be concentrated in areas where the most dynamic changes are occurring.

In this paper, we present a new geodesic image regression framework which decouples the parameterization of the time-varying deformations from the specific representation of the images. This allows the number of model parameters to reflect the complexity of anatomical changes in time rather than the sampling of the images. An upper bound on the dimensionality of the deformation is chosen by the user, and a L^1 penalty selects the most relevant subset of initial momenta, those that describe the most salient changes over time. We apply our sparse geodesic regression model to both synthetic and real data and show that we can decrease the number of model parameters (from the number of voxels down to hundreds) with only minimal decrease in model accuracy.

2 Methods

The goal of image regression is to infer the continuous image evolution which best describes a discrete set of observed images $\mathbf{O}_{t_i}$ at time t_i within the time interval $[t_0,\ T]$. The dynamics of image evolution are modeled as the geodesic flow of diffeomorphisms applied to a baseline image $\mathbf{I}_0$, defined as $\mathbf{I}(t) = \mathbf{I}_0 \circ \phi_t^{-1}$.

Let $\mathbf{y} = (y_1, \ldots, y_M)$ be the physical coordinates of the pixel locations of the observed image furthest in time. The trajectory of the image points can be written as $\mathbf{Y}(t, \mathbf{y}) = \phi_t^{-1}(\mathbf{y})$ where $\mathbf{Y}(0, \mathbf{y}) = \mathbf{y}$ The trajectory $\mathbf{Y}(t, \mathbf{y})$ can then be used to interpolate the gray values of the neighboring voxels in the baseline image, allowing image evolution to be written as $\mathbf{I}(t, \mathbf{y}) = \mathbf{I}_0(\mathbf{Y}(t, \mathbf{y}))$. The variational framework for image regression can be expressed by the criterion

$$\begin{aligned} E(\mathbf{I}_0, \phi_t)) &= \sum_{i=1}^{N_{obs}} ||\mathbf{I}_0(\mathbf{Y}(t_i)) - \mathbf{O}_{t_i})||_{L^2}^2 + \mathrm{Reg}(\phi) \\ &= \sum_{i=1}^{N_{obs}} D(\mathbf{I}_0(\mathbf{Y}(t_i)), \mathbf{O}_{t_i}) + L(\phi) \end{aligned} \tag{1}$$

where D represents the squared distance between images and L is a measure of the regularity of the *entire* time-varying deformation ϕ_t for all times t.

2.1 Discrete parameterization of deformations

The dense flow of diffeomorphisms ϕ_t is constructed by interpolating a discrete set of momenta located at self-interacting control points [3]. Let $\mathbf{c}_0 = \{c_1, ..., c_{N_c}\}$ be a finite set of control points and corresponding initial momenta vectors $\boldsymbol{\alpha}_0 = \{\alpha_1, ...\alpha_{N_c}\}$, which together form the initial state of the system $\mathbf{S}_0 = \{\mathbf{c}_0, \boldsymbol{\alpha}_0\}$.

The set of particles act as initial conditions for the geodesic equations of motion, which define the trajectory of the control points and momenta forward in time:

$$\begin{cases} \dot{c}_i(t) = \sum_{p=1}^{N_c} K(c_i(t), c_p(t))\alpha_p(t) \\ \dot{\alpha}_i(t) = -\sum_{p=1}^{N_c} \alpha_i(t)^t \alpha_p(t) \nabla_1 K(c_i(t), c_p(t)) \end{cases} \tag{2}$$

such that the energy of the system is conserved over time and K is the interpolating kernel governing the interaction between control points. The state of the system $\mathbf{S}(t) = \{c_i(t), \alpha_i(t)\}$ evolves according to these equations, which will be written in short as $\dot{\mathbf{S}}(t) = F(\mathbf{S}(t))$.

The trajectories of control points $c_i(t)$ and $\alpha_i(t)$ parameterize a *time-varying* velocity field

$$v(x,t) = \sum_{p=1}^{N_c} K(x, c_p(t))\alpha_p(t). \tag{3}$$

which is defined at any point in space x and time t. The starting pixel locations $\mathbf{y}$ follow the trajectory $\mathbf{Y}(t, \mathbf{y}) = \phi_t^{-1}(\mathbf{y})$, which evolves in time by the ODE

$$\dot{\mathbf{Y}}(t, \mathbf{y}) = -[d_{\mathbf{y}}\mathbf{Y}(t, \mathbf{y})]v(\mathbf{y}, t) \qquad \text{with } \mathbf{Y}(0, \mathbf{y}) = \mathbf{y}. \tag{4}$$

which is written in short as $\dot{\mathbf{Y}}(t, .) = G(\mathbf{Y}(t, .), \mathbf{S}(t))$.

Using the pixel locations $\mathbf{y}$ as final conditions at time T, integrating this ODE backwards in time computes the inverse deformation of the pixel coordinates from time T to t_0. The deformed pixel locations $\mathbf{Y}(t)$ are used to interpolate the gray values of the neighboring voxels in the baseline image. Therefore the flow of diffeomorphisms is fully determined by the initial state of the system $\mathbf{S}_0$.

2.2 Regression criterion and minimization

The initial state of the system $\mathbf{S}_0 = \{\mathbf{c}_0, \boldsymbol{\alpha}_0\}$ consists of N_c control points and momenta vectors which fully parameterize the geodesic flow of diffeomorphisms ϕ_t in the criterion (2), acting as initial conditions for the flow equations (2). The deformation can then be applied to the pixel coordinates $\mathbf{y}$ according to this flow by solving equation (4) and deformed images can be constructed by interpolating the baseline image $I(t) = \mathbf{I}_0(\mathbf{Y}(t))$. Geodesic image regression is now described by the specific regression criterion

$$E(\mathbf{I}_0, \mathbf{S}_0) = \sum_{i=1}^{N_{obs}} \frac{1}{2\lambda^2} D(\mathbf{Y}(t_i)) + L(\mathbf{S}_0) \tag{5}$$

subject to

$$\begin{cases} \dot{\mathbf{S}}(t) = F(\mathbf{S}(t)) & \text{with } \mathbf{S}(0) = \{\mathbf{c}_0, \boldsymbol{\alpha}_0\} \\ \dot{\mathbf{Y}}(t) = G(\mathbf{Y}(t), \mathbf{S}(t)) & \text{with } \mathbf{Y}(0, \mathbf{y}) = \mathbf{y} \end{cases} \tag{6}$$

where λ^2 is the tradeoff parameter and $L(\mathbf{S}_0) = \sum_{p,q} \alpha^t_{0,p} K(c_{0,p}, c_{0,q}) \alpha_{0,q}$ is the kinetic energy of the particles. The first part of (6) governs the time evolution of the control points and momenta as in (2). The second equation of (6) represents flowing the pixel coordinates along the deformation defined by $\mathbf{S}(t)$ as in (4).

Gradient with respect to control points and initial momenta. The gradient of the criterion (5) with respect to the initial state of the system is

$$\nabla_{\mathbf{S}_0} E = \xi(0) + \nabla_{\mathbf{S}_0} L \tag{7}$$

where the auxiliary variables $\theta(t)$ and $\xi(t) = \{\xi^c,\ \xi^\alpha\}$ satisfy the ODEs:

$$\begin{aligned} \dot{\theta}(t) &= -\partial_1 G(t)^t \theta(t) + \sum_{i=1}^{N_{Obs}} \nabla_{\mathbf{Y}(t_i)} D(t_i) \delta(t - t_i) & \theta(T) = 0 \\ \dot{\xi}(t) &= -(\partial_2 G(t)^t \theta(t) + d_{\mathbf{S}(t)} F(t)^t \xi(t)) & \xi(T) = 0 \end{aligned} \tag{8}$$

The gradient is computed by first integrating equations (2) forward in time to construct the flow of diffeomorphisms. The deformations are then applied by integrating equation (4) backward in time. With the full trajectory of the deformed pixel coordinates, one can compute the gradient of the data term $\nabla_{\mathbf{Y}(t_i)} D(t_i)$, corresponding to each observation. The ODEs (8) are then integrated backwards in time, with the gradients of the data term added at observation time points. The final values of the auxiliary $\xi(t_0)$ is then used to update the location of the control points and initial momenta.

Gradient with respect to baseline image. The gradient with respect to the baseline image $\mathbf{I}_0$ can be computed as the sum of gradients $\nabla_{\mathbf{I}_0} D(\mathbf{Y}(t_i), \mathbf{O}_{t_i})$ for each observed image $\mathbf{O}_{t_i}$. Recall that the intensity of the baseline image $\mathbf{I}_0(Y_k(t))$ for some voxel k and time t is computed by interpolating gray values, written as $\sum_{p \in \mathcal{N}(\mathbf{Y}_k(t))} \rho_p(Y_k(t)) \mathbf{I}_0(\pi_p(Y_k(t)))$ in a neighborhood $\mathcal{N}$ of voxels $\boldsymbol{\pi}$ around $Y_k(t)$ with weights $\boldsymbol{\rho}$ from bilinear (2D) or trilinear (3D) interpolation. Let R_{t_i} be the residual image $\mathbf{I}_0(\mathbf{Y}(t_i)) - \mathbf{O}_{t_i}$. Then a variation $\mathbf{I}_0$ leads to

$$\delta D(\mathbf{Y}(t_i)) = \sum_{i=1}^{N_{obs}} \sum_{k=1}^{N_{vox}} R_{t_i}(Y_k(t_i)) \sum_{p \in \mathcal{N}(\mathbf{Y}_k(t))} \rho_p(Y_k(t)) \delta \mathbf{I}_0(\pi_p(Y_k(t))) \tag{9}$$

$$= \sum_{i=1}^{N_{obs}} \left[\sum_{k=1}^{N_{vox}} \rho_p(Y_k(t_i)) R_{t_i}(Y_k(t_i)) \right] \delta \mathbf{I}_0(\pi_p(Y_i(t_i))). \tag{10}$$

The gradient of D is computed by flowing a voxel $Y_k(t)$ to time t and computing the residual. The gray value in the residual image at voxel k is then distributed to its neighboring voxels with weights defined by bilinear or trilinear interpolation. The summation over observations shows that the gray values are accumulated for every observed image.

Fig. 1. Synthetic image evolution generated by shooting the baseline image (far left) along predefined initial momenta, constraining the resulting evolution to be geodesic.

2.3 Sparsity on initial momenta

To determine the optimal number of control points/momenta, we introduce a L^1 penalty used in the context of atlas building [3] to the regression criterion:

$$E(\mathbf{I}_0, \mathbf{S}_0) = \sum_{i=1}^{N_{obs}} \frac{1}{2\lambda^2} D(\mathbf{Y}(t_i)) + L(\mathbf{S}_0) + \gamma_{\text{sp}} \sum_{i=1}^{N_c} ||\alpha_i(t_0)|| \tag{11}$$

where N_c denotes the total number of control points and α_k^i is the kth initial momentum. We use the Fast Iterative Shrinkage and Thresholding Algorithm [1] to optimize the regression criterion (11) which now contains non-differentiable terms. The update equations for the baseline image, control points, and initial momenta are not affected by the L^1 term. The update equation for momenta can be found in [3], which involves a soft thresholding function to eliminate momenta with small magnitude.

3 Experiments

Synthetic Evolution We first explore a synthetic image time series. The synthetic data, shown in Fig. 1, was generated by shooting the baseline image along predefined initial momenta. The resulting evolution is therefore geodesic, which

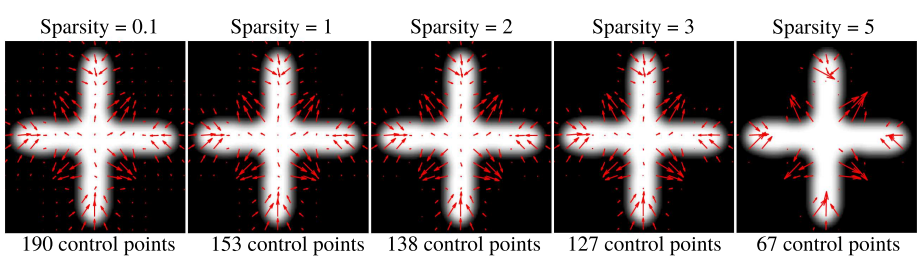

Fig. 2. Baseline shape and initial momenta estimated for several values of sparsity parameter γ_{sp}. An increase in the sparsity parameter leads to a more compact representation, with momenta located in areas of dynamic change over time.

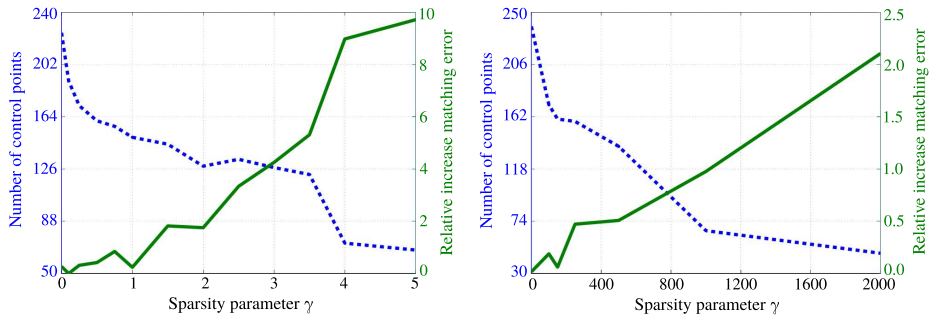

Fig. 3. Impact of the sparsity parameter for the synthetic experiment (left) and for the developing brain (right). There is a range of values of the sparsity parameter which result in a considerable decrease in the number of control points for only minimal increase in the relative error of the data matching term.

serves as a best case scenario for our model. We explore the impact of the sparsity parameter by estimating several geodesic models with a range of values of γ_{sp}. Each model was initialized with control points distributed on a regular grid with 20 pixel spacing, deformation kernel $\sigma_V = 20$ pixels, and $\lambda = 0.5$.

The estimated baseline shape and initial momenta for increasing values of sparsity parameter γ_{sp} are shown in Fig. 2. Increasing the sparsity parameter results in a decrease in the number of control points, leading to a more compact representation of the dynamics of shape change over time. The momenta which remain for high values of γ_{sp} represent the areas that undergo the most dynamic changes. In that sense, these momenta hold the most important information about the trajectory of image evolution. The left panel of Fig. 3 shows the impact the sparsity parameter has on the number of control points as well as the accuracy in which the estimated evolution matches observed data. We can obtain a 43% decrease in the number of control points for less than a 2% relative increase in data matching error, or a 70% decrease in control points for a 10% cost. We obtain a reasonable model of shape change with as few as 67 momenta, compared to 79804 momenta (one for every pixel) for dense deformation models.

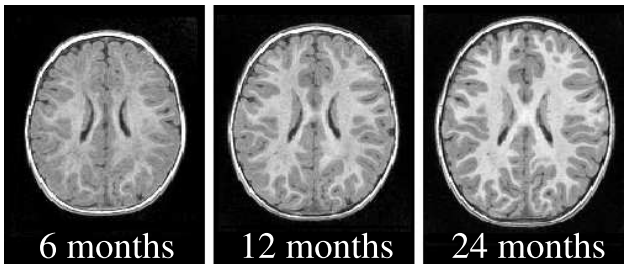

Fig. 4. Observed data acquired from the same child at 6, 12, and 24 months of age.

Pediatric Brain Development Next we investigate the application of sparse geodesic regression to model pediatric brain development. The data considered here, shown in Fig. 4, consists of T1W images acquired from the *same* child at 6, 12, and 24 months old. Due to the limited contrast in the image at 6 months, we estimate geodesic models with baseline image at 24 months, and follow the evolution backwards in time. We estimate several models by varying the sparsity parameter γ_{sp}, with control points initialized on a regular grid with 14 mm spacing, deformation kernel $\sigma_V = 14$ mm, and $\lambda = 0.1$.

Fig. 5 shows the estimated baseline images at 24 months and initial momenta which characterize evolution *backwards* in time. For high values of γ_{sp}, the momenta cluster around the outside of the brain and around the the lateral ventricles. This can be interpreted to mean that the change in brain and ventricle size are the most notable descriptors of this child's trajectory of growth. The right panel of Fig. 3 summarizes the impact of the sparsity parameter in terms of control points and accuracy. We can decrease the number of control points by 80% for just over 2% relative increase in data matching error. This has the potential to improve the power of ensuing statistical analysis, as the discarded momenta can be considered as noise in the description of image evolution.

4 Conclusions

We have presented a sparse geodesic image regression framework that decouples deformation parameters from the specific representation of input images, allowing the dimensionality of the model to be controlled by the user. We further add a L^1 penalty which selects an optimal subset of initial momenta which reflects the complexity of anatomical changes in time instead of the sampling of the image. This reduces the amount of model parameters from the number of image voxels down to tens or hundreds. The reduction of momenta has the

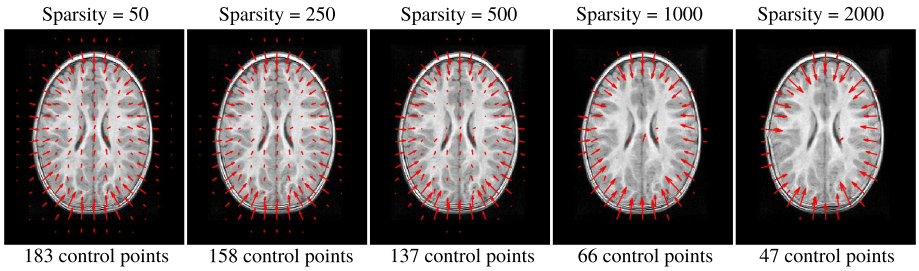

Fig. 5. Baseline shape and initial momenta estimated for several values of sparsity parameter γ_{sp}. The baseline shape was estimated at 24 months of age and evolution was followed backward in time. As the sparsity parameter is increased, the momenta cluster around the perimeter of the brain and around the lateral ventricles. From this, we can infer that the scale of the brain and ventricles are the most salient features describing the development of this child.

potential to improve statistical power, as additional momenta can be considered noise in the description of image change over time. We showed on synthetic and real data that we can greatly reduce the number of momenta, incurring only a minimal cost in terms of matching the target images. Future work will focus on incorporating our sparse geodesic regression model into a framework for longitudinal analysis. Such an integration would require a way to compare sparse sets of momenta across subjects and populations.

References

1. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM J. Img. Sci. 2(1), 183–202 (2009)
2. Davis, B., Fletcher, P., Bullitt, E., Joshi, S.: Population shape regression from random design data. In: ICCV. pp. 1–7. IEEE (2007)
3. Durrleman, S., Allassonnière, S., Joshi, S.: Sparse adaptive parameterization of variability in image ensembles. IJCV pp. 1–23 (2012)
4. Durrleman, S., Pennec, X., Trouvé, A., Braga, J., Gerig, G., Ayache, N.: Toward a comprehensive framework for the spatiotemporal statistical analysis of longitudinal shape data. IJCV pp. 1–38 (2012)
5. Fishbaugh, J., Durrleman, S., Gerig, G.: Estimation of smooth growth trajectories with controlled acceleration from time series shape data. In: Fichtinger, G., Peters, T. (eds.) MICCAI. LNCS, vol. 6892, pp. 401–408. Springer (2011)
6. Fishbaugh, J., Prastawa, M., Durrleman, S., Piven, J., Gerig, G.: Analysis of longitudinal shape variability via subject specific growth modeling. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI. LNCS, vol. 7510, pp. 731–738. Springer (2012)
7. Fletcher, P.: Geodesic Regression on Riemannian Manifolds. In: Pennec, X., Joshi, S., Nielsen, M. (eds.) MICCAI MFCA. pp. 75–86 (2011)
8. Hart, G., Shi, Y., Zhu, H., Sanchez, M., Styner, M., Niethammer, M.: DTI longitudinal atlas construction as an average of growth models. In: Gerig, G., Fletcher, P., Pennec, X. (eds.) MICCAI STIA (2010)
9. Liao, S., Jia, H., Wu, G., Shen, D.: A novel longitudinal atlas construction framework by groupwise registration of subject image sequences. NeuroImage 59(2), 1275–1289 (2012)
10. Lorenzi, M., Ayache, N., Pennec, X.: Schilds ladder for the parallel transport of deformations in time series of images. In: Szekely, G., Hahn, H. (eds.) IPMI. LNCS, vol. 6801, pp. 463–474 (2011)
11. Niethammer, M., Huang, Y., Vialard, F.X.: Geodesic regression for image time-series. In: Fichtinger, G., Peters, T. (eds.) MICCAI. LNCS, vol. 6982, pp. 655–662. Springer (2011)
12. Singh, N., Hinkle, J., Joshi, S., Fletcher, P.: A hierarchical geodesic model for diffeomorphic longitudinal shape analysis. In: IPMI. p. (to appear) (2013)
13. Singh, N., Hinkle, J., Joshi, S., Fletcher, P.: A vector momenta formulation of diffeomorphisms for improved geodesic regression and atlas construction. In: ISBI. p. (to appear) (2013)
14. Vialard, F., Trouvé, A.: Shape splines and stochastic shape evolutions: A second-order point of view. Quarterly of Applied Mathematics 70, 219–251 (2012)