

Estimation of Smooth Growth Trajectories with Controlled Acceleration from Time Series Shape Data

James Fishbaugh, Stanley Durrleman, and Guido Gerig

Scientific Computing and Imaging Institute
University of Utah
Salt Lake City, Utah

Abstract. Longitudinal shape analysis often relies on the estimation of a realistic continuous growth scenario from data sparsely distributed in time. In this paper, we propose a new type of growth model parameterized by acceleration, whereas standard methods typically control the velocity. This mimics the behavior of biological tissue as a mechanical system driven by external forces. The growth trajectories are estimated as smooth flows of deformations, which are twice differentiable. This differs from piecewise geodesic regression, for which the velocity may be discontinuous. We evaluate our approach on a set of anatomical structures of the *same* subject, scanned 16 times between 4 and 8 years of age. We show our acceleration based method estimates smooth growth, demonstrating improved regularity compared to piecewise geodesic regression. Leave-several-out experiments show that our method is robust to missing observations, as well as being less sensitive to noise, and is therefore more likely to capture the underlying biological growth.

1 Introduction

The study of time dependent shapes is an emerging field in Computational Anatomy, with potential applications to early brain development, aging studies, or the analysis of evolving pathologic structures. As longitudinal data becomes more widely available, the need for computer models of anatomical evolution becomes increasingly important. Two approaches have been followed so far: the first consists in computing a realistic growth scenario from cross-sectional time-series data, like in [4,10,6,3]. The second approach involves estimating several individual growth trajectories and combining them with a framework for 4D registration between growth trajectories or 4D atlas construction, to statistically analyze the growth variability within a population, like in [13,8,14,7,9].

In any case, the methods rely greatly on the estimation of growth models from time series data, which are sparsely distributed in time. Growth models provide a tool to generate shapes at any instant in time (within the interval defined by the data), offering us the opportunity to continuously measure shape properties. This is in contrast to using sparse measurements such as volume or circumference

for 1D regression absent the shape information. The problem can be stated as “temporal shape regression” and can be solved by purely descriptive statistical methods like the extension of kernel regression to Riemannian manifolds [4], or by generative statistical models which define a parameterized family of realistic growth models and the one which best fits the actual data is estimated based on a regularized least-square criterion [8,7,2]. We favor this last approach, since it makes explicit the assumptions which drive the estimation of growth trajectories and therefore enables the inclusion of realistic biological priors to constrain the estimation.

The growth model in [7] is based on a continuous flow of diffeomorphisms, with piecewise geodesics interpolating between shapes. This method estimates continuous non-linear growth between shapes, but does not guarantee differentiable growth as the speed of evolution is discontinuous at observation time-points. Our work is motivated by the assumption that the evolution of biological tissue is inherently *smooth* in time. If we consider the growth of biological tissue as a mechanical system driven by external forces, then the evolution of any particle on an anatomical surface is continuous with continuous derivative and therefore does not change direction *instantaneously*, as observed in the growth model estimated from [7].

Temporal smoothness can be enforced via smooth interpolation between 3D deformations estimated at discrete time-points, using B-splines or polynomial interpolation as in [11,1]. However, these approaches are not based on the inference of a generic growth model, which captures the dynamics of the shape changes over time.

Based on these considerations, we propose a new growth model parameterized by acceleration, rather than velocity as in the large deformation setting of [12]. The estimated acceleration could be considered an indication of the forces which drive the growth of the anatomical structures. From this parameterization, we gain one order of differentiability and guarantee that shape evolution is smooth in both space and time. We further deviate from the large deformations framework by introducing a new regularization term which accounts for the total amount of acceleration. As a consequence, our model does not constrain the flow of deformation between shapes to be geodesic, or close to a geodesic path. By contrast, the approach in [16] estimates twice differentiable trajectories as random perturbations of geodesic paths.

The evaluation of our new methodology on real anatomical surfaces reveals the differences between our approach and piecewise geodesic regression. Our regression yields a twice differentiable evolution with improved regularity, thus discarding more noise from the data to fit a more realistic growth trajectory. Also, we demonstrate that volume measurements taken out of our 3D shape regression are compatible with a 1D regression of these measurements, whereas piecewise geodesic regression appears to overfit. Lastly, we show via leave-several-out experiments that our method better interpolates between data and is therefore more robust to missing observations. This suggests a greater ability to capture the underlying growth of the anatomical structures.

2 Shape Regression Parameterized by Acceleration

The problem of longitudinal shape regression involves inferring a continuous shape evolution from a discrete set of shapes S_{t_i} observed at time t_i . Shape evolution is modeled as the continuous deformation of a baseline shape S_0 , formally defined as $R_t = \phi_t(S_0)$ where R_t corresponds to S_0 having undergone the transformation ϕ_t with t varying continuously within the time interval. The time-varying deformation ϕ_t is a general transformation from \mathbb{R}^N to \mathbb{R}^N with $\phi_0(S_0) = S_0$. The baseline shape is deformed over time to closely match the observed shapes ($R_{t_i} \sim S_{t_i}$) while the rigidity of the deformation is controlled via a regularity term. This leads to a variational problem in the form of a trade off between fidelity to data and regularity. For measuring shape similarity, we follow the work of [15], modeling shapes as currents.

We define the acceleration field $a(x, t)$ at point x and time t as a vector field of the form

$$a(x, t) = \sum_{i=1}^N K^V(x, x_i(t))\alpha_i(t) \tag{1}$$

where x_i are the shape points carrying a point force vector α_i , and $K^V(x, y) = \exp(-\|x - y\|^2 / \lambda_V^2)$ is a Gaussian kernel of dimension mass^{-1} with standard deviation λ_V controlling the spatial extent at which the acceleration field varies.

The time-varying point force vectors $\alpha_i(t)$ parameterize a flow of deformation $\phi_t(x_i(t))$ by the integration of the 2nd-order ODE $\ddot{\phi}_t(x_i(t)) = a(x_i(t), t)$ with initial position $x_i(0)$ and initial velocity $\dot{x}_i(0)$. The initial positions of the particle are assumed to be fixed at the vertices of the baseline shape, while the initial velocity of the particles have to be determined by the algorithm.

Let $\mathbf{x}(t)$, $\mathbf{a}(t)$, and $\boldsymbol{\alpha}(t)$ be the concatenation of the $x_i(t)$'s, $a_i(t)$'s, and the $\alpha_i(t)$'s. This parameterization leads to the specific regression criterion

$$E(\dot{\mathbf{x}}(0), \boldsymbol{\alpha}(t)) = \sum_{t_i} \|\phi_{t_i}(\mathbf{x}(0)) - \mathbf{x}(t_i)\|_{W^*}^2 + \gamma \int_0^T \|\mathbf{a}(t)\|_V^2 dt \tag{2}$$

where $\|\cdot\|_{W^*}$ is the norm on currents and regularity is defined as $\|\mathbf{a}(t)\|_V^2 = \boldsymbol{\alpha}(t)K^V(\mathbf{x}(t), \mathbf{x}(t))\boldsymbol{\alpha}(t)$, interpreted as the ‘total amount of acceleration’, measured using the norm in the reproducing kernel Hilbert space defined by the interpolating kernel [5].

3 Description of the Algorithm

We implement an adaptive step size gradient descent algorithm. The gradient of the criterion (2) with respect to force vectors and initial velocity is written as

$$\nabla_{\alpha_{i(t)}} E(t) = 2\gamma\alpha_i(t) + \eta_i^{\dot{x}}(t) \quad \text{and} \quad \nabla_{\dot{x}_i(0)} E = \eta_i^{\dot{x}}(0) \tag{3}$$

where variables $\eta_i^x(t)$ and $\eta_i^{\dot{x}}(t)$ satisfy coupled ODEs shown in Appendix A.

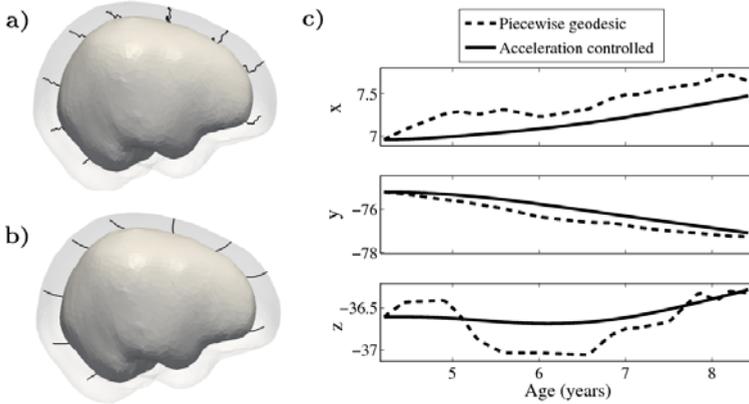


Fig. 1. **a)** and **b)** Shape evolution from baseline (solid) to final configuration (transparent) using a model based on piecewise geodesics (a) and our method (b) with point trajectories for selected particles displayed as black lines. **c)** The path of a point on the forebrain is decomposed into coordinates. Growth is estimated using 15 target shapes, highlighting the speed discontinuities present in the piecewise geodesic evolution.

During each iteration of gradient descent, the trajectories of shape points are computed by solving the 2nd-order ODE $\ddot{\phi}_t(x_i(t)) = a(x_i(t), t)$ using a Verlet integration scheme. The auxiliary variables $\eta_i^x(t)$ and $\eta_i^{\dot{x}}(t)$ are computed using an Euler method with prediction/correction. Eventually we compute the gradients given in equation (3). The algorithm may be started with zero initial velocity and force, though we notice faster convergence when initial velocity is determined by geodesic diffeomorphic registration between the baseline and first target shape as in [15].

4 Experiments

To evaluate our method, we use longitudinal image data from a child that has been scanned 16 times between four and eight years of age. The MRI data is first rigidly aligned to establish a common reference frame. The intracranial volume and lateral ventricles are segmented from each image using an EM based tissue classification algorithm and a level-set based active contour segmentation tool.

We estimate the evolution of the intracranial surface using a regression model based on the piecewise geodesic flow of diffeomorphisms as in [7]. The standard deviation of the Gaussian kernel controlling deformation is set to 50 mm, roughly 30% of the diameter of the baseline intracranial surface. For the scale of currents we use 20 mm, with a regularity weight of 0.1. Finally, time is discretized in increments of 0.0425 years. We also produce a growth trajectory using our proposed method with the same parameter settings as above except we weight regularity by 0.01 (the two weighted terms cannot be compared since they have different ‘physical’ dimension). The parameters were tuned empirically to produce regressions of comparable quality with both methods.

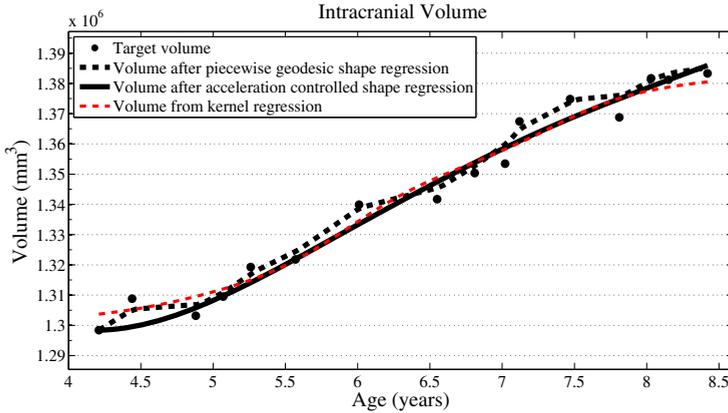


Fig. 2. Volume measurements derived from our growth model are consistent with a kernel regression ($\sigma = 0.5$) performed on the sparse volume measurements. Our model describes the continuous evolution of *shape* and volume is measured after regression.

Shape evolution is considerably smoother using our proposed regression model as compared to the piecewise geodesic model. This is particularly evident in the trajectories of the shape points across time, a subset of which are shown in Fig. 1. It is an important distinction that the trajectories estimated by our method are *not* a smoothing of the piecewise geodesic method. Rather, the trajectories are the result of fundamentally different assumptions on the underlying model which results in a more realistic estimation of growth.

The smoothness constraints imposed by our model limit the shape variation we can capture over short time periods. Consequently, we investigate the accuracy of our model by examining how closely we match the target data: our estimated growth scenario decreases the initial sum of squared residual by 148%, compared to a 153% decrease from the piecewise geodesic method. While our method does not come as close to matching the target data, this suggests that our method is less sensitive to noise and less likely to overfit.

Next, we investigate the application of our model to the study of measurements derived from shape. Here we obtain a continuous non-linear model of volume, shown in Fig. 2. The results are consistent with a 1D regression model, such as kernel regression, applied to the sparse volume measurements. However, we have focused our modeling efforts on capturing the evolution of shape, with continuous volume measurements resulting naturally from the estimated growth. In addition, the piecewise geodesic method appears to be overfitting, producing unrealistic volume measurements, further suggesting that our method is more robust in the presence of noisy data.

Finally, we consider the evolution of the lateral ventricles, which exhibit considerably more complexity than the intracranial surface. The horns of the segmented lateral ventricles are as thin as a few millimeters, making regression particularly challenging. As with the intracranial volume, ventricle growth is estimated using a piecewise geodesic model and our acceleration based model. The

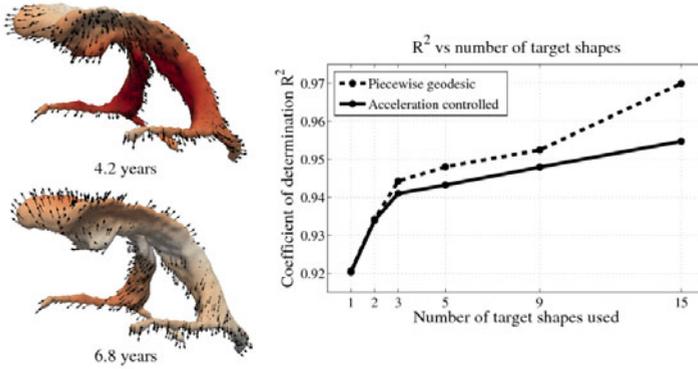


Fig. 3. Left: Snapshots from a continuous shape evolution of lateral ventricles estimated by our regression model. Acceleration vectors are displayed on the surface, with color denoting magnitude. **Right:** The impact of the number of target shapes on R^2 .

scale of deformation is set at 6 mm, the scale of currents to 2 mm, and regularity is weighted by 0.1 and 0.01, respectively.

The impact of missing data is examined by performing leave-several-out experiments, the results of which are summarized in Fig. 3. In all experiments, selected target shapes were chosen as uniformly across time as possible. Our method demonstrates robustness with respect to the number of target shapes, with only minimal increase in the coefficient of determination R^2 when using more than 3 targets. This suggests that our method captures the underlying growth with limited data, as additional target data does not greatly alter the estimation. In contrast, piecewise geodesic regression is more influenced by additional target data and is therefore likely to overfit.

5 Conclusion

We have introduced a new 2nd-order regression model for estimating smooth evolution from a collection of time dependent shape data. This is based on a new way of parameterizing growth based on acceleration rather than velocity. We show on real anatomical data that, compared to the standard piecewise geodesic model, our method is less sensitive to noise introduced during segmentation and is robust to missing data, and is therefore more likely to characterize the underlying biological growth. The evolution of volume extracted after shape regression was shown to be compatible with a 1D regression on the observed volume measurements. Our method may be improved by additionally solving for initial positions of the shape points as in [7], to address the apparent underestimation of initial volume in Fig. 2.

Note that the new concept introduced in this paper has been implemented for 3D-surface data modeled as currents but can be easily adapted to a variety of other data and metrics. Future work will focus on the interpretation of the estimated acceleration in terms of external forces exerted on the biological tissue. This will enable the addition of more biological and mechanical priors.

Acknowledgments. This work was supported by NIH grant RO1 HD055741 (ACE, project IBIS) and by NIH grant U54 EB005149 (NA-MIC).

References

1. Castillo, E., Castillo, R., Martinez, J., Shenoy, M., Guerrero, T.: Four-dimensional deformable image registration using trajectory modeling. *Physics in Medicine and Biology* 55, 305–327 (2010)
2. Craene, M.D., Camara, O., Bijmens, B.H., Frangi, A.F.: Large diffeomorphic FFD registration for motion and strain quantification from 3D-US sequences. In: Ayache, N., Delingette, H., Sermesant, M. (eds.) FIMH 2009. LNCS, vol. 5528, pp. 437–446. Springer, Heidelberg (2009)
3. Datar, M., Cates, J., Fletcher, P., Gouttard, S., Gerig, G., Whitaker, R.: Particle based shape regression of open surfaces with applications to developmental neuroimaging. In: Yang, G.Z., Hawkes, D. J., Rueckert, D., Noble, J.A., Taylor, C. J. (eds.) MICCAI 2009. LNCS, vol. 5762, pp. 167–174. Springer, Heidelberg (2009)
4. Davis, B., Fletcher, P., Bullitt, E., Joshi, S.: Population shape regression from random design data. In: Proc. of ICCV, pp. 1–7 (October 2007)
5. Durrleman, S.: Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution. Thèse de sciences (phd thesis). Université de Nice-Sophia Antipolis (March 2010)
6. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N., Braga, J.: Comparison of the endocast growth of chimpanzees and bonobos via temporal regression and spatiotemporal registration. In: MICCAI-STIA Workshop (2010)
7. 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)
8. Ehrhardt, J., Werner, R., Schmidt-Richberg, A., Schulz, B., Handels, H.: Generation of a mean motion model of the lung using 4D-CT image data. In: Proc. of Eurographics Workshop on VCBM, pp. 69–76 (2008)
9. Hart, G., Shi, Y., Zhu, H., Sanchez, M., Styner, M., Niethammer, M.: DTI longitudinal atlas construction as an average of growth models. In: MICCAI-STIA Workshop (2010)
10. Khan, A., Beg, M.: Representation of time-varying shapes in the large deformation diffeomorphic framework. In: Proc. of ISBI, pp. 1521–1524 (2008)
11. Metz, C., Klein, S., Schaap, M., van Walsum, T., Niessen, W.: Nonrigid registration of dynamic medical imaging data using nd+t b-splines and a groupwise optimization approach. *Medical Image Analysis* 15(2), 238–249 (2011)
12. Miller, M.I, Trouvé, I., Younes, L.: On the metrics and Euler-Lagrange equations of Computational Anatomy. *Annual Review of Biomedical Engineering* 4, 375–405 (2002)
13. Perperidis, D., Mohiaddin, R.H., Rueckert, D.: Spatio-temporal free-form registration of cardiac MRI sequences. *Medical Image Analysis* 9(5), 441–456 (2005)
14. Peyrat, J.M., Delingette, H., Sermesant, M., Pennec, X., Xu, C., Ayache, N.: Registration of 4D Time-Series of Cardiac Images with Multichannel Diffeomorphic Demons. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 972–979. Springer, Heidelberg (2008)

- 15. Vaillant, M., Glaunès, J.: Surface matching via currents. In: Christensen, G.E., Sonka, M. (eds.) IPMI 2005. LNCS, vol. 3565, pp. 381–392. Springer, Heidelberg (2005)
- 16. Vialard, F., Trounev, A.: Shape splines and stochastic shape evolutions: A second-order point of view. Quarterly of Applied Mathematics To Appear

A Differentiation of the Regression Criterion

Using matrix notation, we denote the current state of the system of shape points by the vector $\mathbf{X}(t) = (\mathbf{x}(t), \dot{\mathbf{x}}(t))^t$ concatenating position and velocity of every point. The state of the system is evolved by the following differential equation:

$$\dot{\mathbf{X}}(t) = F(\mathbf{X}(t), \boldsymbol{\alpha}(t)) = \begin{pmatrix} \dot{\mathbf{x}}(t) \\ \ddot{\mathbf{x}}(t) = \mathbf{K}(\mathbf{x}(t), \mathbf{x}(t))\boldsymbol{\alpha}(t) \end{pmatrix} \tag{4}$$

with initial condition $\mathbf{X}(0) = \mathbf{X}_0 = (\mathbf{x}_0, \dot{\mathbf{x}}_0)^t$.

We now rewrite (2) as $E(\mathbf{X}(t)) = \sum_{t_i} A(\mathbf{X}(t_i)) + \gamma \int_0^T L(\mathbf{X}(t), \boldsymbol{\alpha}(t))dt$. Let δE be a variation of the criterion E with respect to a variation $\delta\boldsymbol{\alpha}(t)$ of the impulse vectors $\boldsymbol{\alpha}(t)$, which induces a variation of the state variable $\mathbf{X}(t)$:

$$\delta E = \sum_{t_i} (d_{\mathbf{X}(t_i)} A_i) \delta\mathbf{X}(t_i) + \gamma \int_0^T (\partial_{\mathbf{X}(t)} L(t))\delta\mathbf{X}(t) + (\partial_{\boldsymbol{\alpha}(t)} L(t))\delta\boldsymbol{\alpha}(t)dt \tag{5}$$

The ODE in (4) shows that these variations $\delta\mathbf{X}(t)$ satisfy a linear inhomogeneous ODE. The method of variation of parameters gives the solution

$$\delta\mathbf{X}(t) = R_{0t}\delta\mathbf{X}_0 + \int_0^T R_{ut}\partial_{\boldsymbol{\alpha}(u)} F(u)\delta\boldsymbol{\alpha}(u)\mathbf{1}_{\{u \leq t_i\}}du \tag{6}$$

where $R_{ut} = \exp\left(\int_u^t \partial_{\mathbf{X}(s)} F(s)ds\right)$ and $\mathbf{1}_{\{t \leq t_i\}} = 1$ if $t \leq t_i$ and 0 otherwise.

Plugging this equation into (5) leads to:

$$\nabla_{\boldsymbol{\alpha}} E(t) = \partial_{\boldsymbol{\alpha}(t)} L(t)^t + \partial_{\boldsymbol{\alpha}(t)} F(t)^t \boldsymbol{\eta}(t) \quad \text{and} \quad \nabla_{\mathbf{X}_0} E = \boldsymbol{\eta}(0) \tag{7}$$

where we denote the auxiliary variable $\boldsymbol{\eta}(t)$ as

$$\boldsymbol{\eta}(t) = \sum_i \nabla_{\mathbf{x}_{t_i}} A_i \mathbf{1}_{\{t \leq t_i\}} + \int_t^T \partial_{\mathbf{X}} L(u)^t + \partial_{\mathbf{X}} F(u)^t \boldsymbol{\eta}(u)du \tag{8}$$

From now on, we decompose the vectors into 2 blocks (the x -component and the \dot{x} -component). Due to the definition of A , L and F , we have $\nabla_{\mathbf{X}(t_i)} A_i = (\nabla_{\mathbf{x}_i} A_i \quad \mathbf{0})^t$, $\partial_{\mathbf{X}} L = (\gamma\boldsymbol{\alpha}^t(\partial_1 + \partial_1)(\mathbf{K}(\mathbf{x}, \mathbf{x})\boldsymbol{\alpha}) \quad \mathbf{0})^t$, $\partial_{\boldsymbol{\alpha}} L = 2\gamma\boldsymbol{\alpha}^t \mathbf{K}(\mathbf{x}, \mathbf{x})$, $\partial_{\mathbf{X}} F = \begin{pmatrix} \mathbf{0} & \mathbf{1} \\ (\partial_1 + \partial_2)\mathbf{K}(\mathbf{x}, \mathbf{x})\boldsymbol{\alpha} & \mathbf{0} \end{pmatrix}$ and $\partial_{\boldsymbol{\alpha}} F^x = (\mathbf{0} \quad \mathbf{K}(\mathbf{x}, \mathbf{x}))$.

Therefore, the gradient of the regression criterion with respect to the L^2 metric given in (7) is now equal to: $\nabla_{\boldsymbol{\alpha}} E(t) = \mathbf{K}(\mathbf{x}(t), \mathbf{x}(t)) (2\gamma\boldsymbol{\alpha}(t) + \boldsymbol{\eta}^{\dot{x}}(t))$, where we have decomposed the auxiliary variable $\boldsymbol{\eta}$ into $\boldsymbol{\eta} = (\boldsymbol{\eta}^x, \boldsymbol{\eta}^{\dot{x}})$.

The matrix $\mathbf{K}(\mathbf{x}(t), \mathbf{x}(t))$ is precisely the Sobolev metric induced by the kernel on the set of L^2 functions, so the gradient is given in coordinates as in (3).