Groupwise and Local Analysis of Image Warps

Carole Twining and Stephen Marsland

Division of Imaging Science and Biomedical Engineering
University of Manchester

1 Introduction

Most non-rigid image registration algorithms are based on aligning pairs of images, but there has been recent interest in groupwise algorithms, which enable the registration of a set of images into a single common frame of reference. The set of deformation fields obtained from such a registration contains information about the variability of structures across the group, meaning that the quantitative analysis and modelling of the set of deformation fields provides useful information about the images in the set. Such analysis must be based (either implicitly or explicitly) on a particular *common* mathematical representation of the set of deformation fields.

Previous work on the analysis of shape variability has used a range of representations; examples include (Cootes et al., 1995; Bookstein, 1989; Pizer et al., 1999). Recent work on modelling dense 2D and 3D deformation fields has either used the densely-sampled deformation vectors directly (e.g., LeBriquer and Gee (1997)), or has employed a smooth, continuous representation of them (e.g.,B-splines, Rueckert et al. (2001)).

In this paper we describe a class of representations for such deformation fields based on *interpolating* splines. We present both diffeomorphic and non-diffeomorphic representations; the non-diffeomorphic is computationally much simpler, and suitable for use within a non-rigid registration algorithm (Marsland and Twining, 2003), whereas the diffeomorphic representation leads to useful metrics for the analysis of a set of parameterised warps. We demonstrate that these metrics are superior to ad hoc Euclidean metrics for the task of classifying legal and illegal variations of a set of shapes. We also discuss differential-geometric constructions of invariant scalars for the analysis of the spatial variation of a warp.

2 Representing Deformation Fields

2.1 Interpolating Splines

Suppose that we have a vector-valued deformation field $\vec{v}(\vec{x})$, and that we know the deformation at some set of knotpoints $\{\vec{x}^i : i = 1 \text{ to } N\}$, where $\vec{v}(\vec{x}^i) = \vec{v}^i$. We say that the spline interpolant for this deformation field is the minimiser of a functional Lagrangian of the form:

$$E[\vec{v}] = \int_{\mathbb{R}^n} d\vec{x}\, |L\vec{v}(\vec{x})|^2 + \sum_{i=1}^{N} \lambda_i \left(\vec{v}(\vec{x}^i) - \vec{v}^i\right) = \|\vec{v}(\vec{x})\|_L^2 + \sum_{i=1}^{N} \lambda_i \left(\vec{v}(\vec{x}^i) - \vec{v}^i\right), \qquad (1)$$

where L is some scalar differential operator, with associated functional norm $\|\cdot\|_L$. The first term in the Lagrangian is the smoothing or energy term; the second term with the Lagrange multipliers $\{\lambda_i\}$ ensures that the spline fits the data at the knotpoints. The choice of operator L and boundary conditions defines a particular spline basis. The general solution can be written:

$$\vec{v}(\vec{x}) = \vec{g}(\vec{x}) + \sum_{i=1}^{N} \vec{\alpha}^i G(\vec{x}, \vec{x}^i), \tag{2}$$

where:

$$\text{the affine function } g \text{ is a solution of: } L\vec{g}(\vec{x}) = 0, \tag{3}$$

$$\text{the Green's function } G \text{ is a solution of: } \left(L^\dagger L\right) G(\vec{x}, \vec{y}) \propto \delta(\vec{x} - \vec{y}), \tag{4}$$

with $L^\dagger$ being the Lagrange dual of L, and $\{\vec{\alpha}^i\}$ the spline coefficients. The table shows some commonly-used Green's functions, $\mathbb{D}^n$ denotes the unit ball in $\mathbb{R}^n$, and $\partial\mathbb{D}^n$ is its boundary.

Name	Dim	$L^\dagger L$	Boundary conditions on $f(\vec{x})$	$G(\vec{x}, \vec{y})$
thin-plate (Duchon, 1976)	even	$(\nabla^2)^2$	asymptotically linear	$\lvert\vec{x} - \vec{y}\rvert^{4-n} \log \lvert\vec{x} - \vec{y}\rvert$
thin-plate (Duchon, 1976)	odd	$(\nabla^2)^2$	asymptotically linear	$\lvert\vec{x} - \vec{y}\rvert^{4-n}$
biharmonic clamped-plate (Boggio, 1905)	2	$(\nabla^2)^2$	$f = f' = 0$ on and outside $\partial\mathbb{D}^2$	$\lvert\vec{x} - \vec{y}\rvert^2 (A^2 - 1 - \log A^2)$, $A(\vec{x}, \vec{y}) = \frac{\sqrt{\vec{x}^2\vec{y}^2 - 2\vec{x}\cdot\vec{y} + 1}}{\lvert\vec{x} - \vec{y}\rvert}$
triharmonic clamped-plate (Boggio, 1905)	3	$(\nabla^2)^3$	$f = f' = 0$ on and outside $\partial\mathbb{D}^3$	$\lvert\vec{x} - \vec{y}\rvert \left(A + \frac{1}{A} - 2\right)$, $A(\vec{x}, \vec{y}) = \frac{\sqrt{\vec{x}^2\vec{y}^2 - 2\vec{x}\cdot\vec{y} + 1}}{\lvert\vec{x} - \vec{y}\rvert}$
Gaussian	n	$e^{-\left(\frac{\nabla^2}{4\beta}\right)}$	asymptotically linear	$\exp\left(-\beta \lvert\vec{x} - \vec{y}\rvert^2\right)$

We can now view a set of knotpoints and knotpoint displacements as a representation of a general displacement field in a given spline basis. A particular spline basis is defined by specifying the form of the differential operator L and the boundary conditions on the spline. However, we note that these displacement fields are only guaranteed to be *diffeomorphic* in the limit of small displacements. Representing diffeomorphic displacement fields, and hence the diffeomorphism group, is dealt with in the next section.

2.2 Geodesic Interpolating Splines

The usual approach to constructing a large-deformation diffeomorphism is to consider such a deformation as an infinite sequence of infinitesimal deformations (Dupuis et al., 1998; Joshi and Miller, 2000). Such a situation can be represented using a time-varying deformation field (flowfield) $\vec{v}(\vec{x}, t)$, where t is the flow time, with an associated energy (the generalisation of equation (1)):

$$E_{opt} = \min_{\vec{v}} \left[\int_0^1 dt \ \|\vec{v}(\vec{x}, t)\|_L^2 \ + \ \sum_{i=1}^{N} \lambda_i \int_0^1 dt \left| \vec{v}(\vec{x}^i(t), t) - \frac{d\vec{x}^i(t)}{dt} \right|^2 \right]. \tag{5}$$

It was shown by Dupuis et al. (1998) that such a flowfield always generates a diffeomorphism if the functional norm is finite; note that this imposes smoothness constraints on the flowfield in space but not in time. As before, the deformation field $\vec{v}(\vec{x}, t)$ (for fixed t) is represented by the positions of some set of knotpoints $\{\vec{x}^i(t)\}$; the time-varying deformation field is then represented by the paths of these knotpoints, which need not be smooth in time (see above). Note that we no longer have an exact solution for the functional minimizer, since the knotpoint paths are only constrained at their end-points. We therefore have to numerically optimise the expression for the energy in equation (5). We use the optimisation scheme previously described in Twining et al. (2002).

This enables us to represent a general element of the diffeomorphism group in question (defined by the particular boundary conditions) in terms of the displacements of a set of knotpoints. For such a set of knotpoints, there will be some optimal flowpath for the knotpoints, with minimal energy E_{opt}. We now look at the properties of this energy, regarded as a function on the group of diffeomorphisms.

3 Going Metric

3.1 The Metric on the Diffeomorphism Group

As in the previous section, we consider a time-varying deformation field $\vec{v}(\vec{x},t)$ and its associated diffeomorphism. Any such flowfield has an associated energy:

$$E[\vec{v}] = \int_0^1 dt \ \|\vec{v}(\cdot,t)\|_L^2 \,. \tag{6}$$

For a given element g of the diffeomorphism group there will be an optimal flowfield $\vec{v}^g(\vec{x},t)$ that generates the given diffeomorphism with the minimum energy, $E_g = E[\vec{v}^g]$. We can then define the function:

$$d(e,g) = d_g = \sqrt{E_g}, \ \text{ where } e \text{ is the identity element of the group.} \tag{7}$$

3.1.1 The Triangle Inequality

Let f and g be two arbitrary elements of the diffeomorphism group, with associated (optimum) flowfields $\vec{v}^f$ and $\vec{v}^g$ and associated (optimum) energies $E_f = d_f^2$, $E_g = d_g^2$. Consider the composition of these 2 group elements $h = f \circ g$, where we define the composition such that $h(\vec{x}) \equiv f(g(\vec{x}))$, and let h have optimum energy E_h. By definition, we can reach h using the following composition of flowfields:

$$\vec{V}(\cdot,t) = \frac{1}{t'}\vec{v}^g\left(\cdot,\frac{t}{t'}\right) \text{ if } 0 \le t \le t' \text{ and } \vec{V}(\cdot,t) = \frac{1}{1-t'}\vec{v}^f\left(\cdot,\frac{t-t'}{1-t'}\right) \text{ if } t' < t \le 1. \tag{8}$$

It is trivial to show that:

$$E[\vec{V}] = \frac{E_g}{t'} + \frac{E_f}{1-t'}, \ \min_{t'} E[\vec{V}] = \left(\sqrt{E_f}+\sqrt{E_g}\right)^2 = (d_f+d_g)^2 \,. \tag{9}$$

We know that $E_h = d_h^2$ is optimal, which gives the inequality:

$$d_h = d_{f\circ g} \le d_f + d_g \Rightarrow d(e, f\circ g) \le d(e,f) + d(e,g). \tag{10}$$

In equation (7) we only defined the function $d(e,\cdot)$. We now extend this definition so that $d(g, f\circ g) \equiv d(e,f)$ for all f and g. (Note that this then makes d right-invariant, whereas Camion and Younes (2001) used a left-invariant formulation; this is because of our different definition of group composition). We can also see that d has to be a symmetric function; reversing the path means changing the sign of the flowfield, which does not effect the energy. Hence, from equation (10):

$$d(h, f\circ g\circ h) \le d(h, g\circ h) + d(g\circ h, f\circ g\circ h) \ \Rightarrow d(h,k) \le d(h,l) + d(l,k) \,\forall\, h,k,l, \tag{11}$$

which is the triangle inequality. Therefore, $d(\cdot,\cdot)$ is a distance function on the group of diffeomorphisms. We cannot extract an explicit form for the local metric tensor since the space is infinite-dimensional.

3.2 The Induced Metric on the Space of Knotpoint Positions

We can consider the metric on the space of knotpoints induced by the distance function on the diffeomorphism group – the mapping between the spaces is constructed by calculating the GIS warp given by the knotpoint positions. So, for knotpoint positions $\{\vec{x}^i \equiv \{x^i_\mu\}\}$, the induced metric tensor has the explicit form:

$$g_{\alpha\beta} = \left(G^{-1}_{ij}\delta_{\mu\nu}\right) \quad \text{where } G_{ij} \equiv G(\vec{x}^i, \vec{x}^j), \tag{12}$$

using the multi-index notation $\alpha = (i, \mu), \beta = (j, \nu)$. This then corresponds to the metric on the space of knotpoint positions given by Camion and Younes (2001). However, in their derivation, they only considered inexact matching between the knotpoint paths and the flowfield, whereas we have imposed exact matching. Also, the derivation given above emphasises the fact that we can directly compare geodesic distances calculated on spaces with varying numbers of knotpoints, since they are all induced from the same metric on the diffeomorphism group.

4 Undergoing Analysis

4.1 Global Analysis

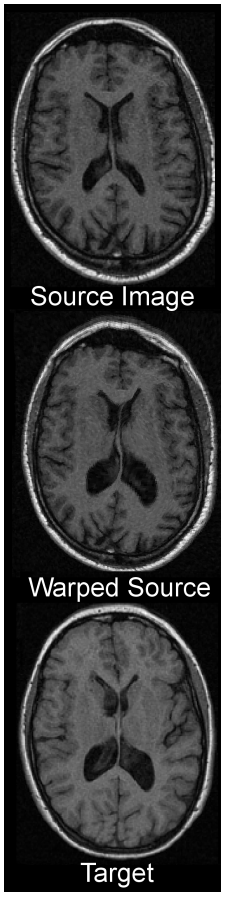

Figure 1:

To study the utility of the metric on the diffeomorphism group, we took a dataset that consisted of a set of 2D T1-weighted MR axial slices of brains, where the slices were chosen to show the lateral ventricles. For each image, the positions of the lateral ventricles and the skull were annotated by a radiologist using a set of 163 points. We took a subset of 66 of these points to be the positions of our knotpoints. Given a pair of images, the knotpoint positions on the images give us the initial and final positions of our knotpoint paths. We then calculated the GIS warp corresponding to these positions using the 2D clamped-plate spline as Green's function. An example warp is shown in figure 1.

We then took the annotated outlines of the anterior lateral ventricles, each consisting of 40 knotpoints. A linear SSM was built from this training set of shapes. We then used this SSM to generate random example shapes. These examples were classified as legal if the outlines of both ventricles did not intersect either themselves or each other, and illegal otherwise – examples of both are shown in figure 2. The training set of shapes are, by definition, legal. We then calculated the GIS warps between the classified set of example shapes and the mean shape from the model. The geodesic distance for the warp from the mean is compared with the Mahalanobis distance in figure 3. It is immediately apparent that we cannot separate the legal and illegal shapes by using the Mahalanobis distance. However, using the geodesic distance, it is possible to construct a simple classifier (shown by the dotted grey line) that separates the two groups, with only one example shape being misclassified (the grey circle just below the line). Given that the Mahalanobis distance for the SSM is equivalent to a Euclidean metric on the space of point deformations, this demonstrates the superiority of the GIS metric over an ad hoc metric. See Twining and Marsland (2003) for further details.

We note that the metric can also be used to calculate the distances between any pair of warps and the geodesic between them, which means that the mean of any pair of warps can be computed. The curvature of the space at any point can also be computed using the Ricci curvature. It is still an unsolved problem to construct generative models on this curved space.

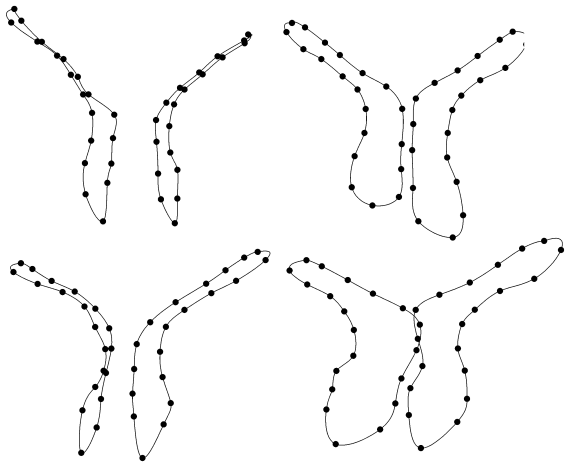

Figure 2: *Top:* Examples from the training set. *Bottom:* Legal (left) and illegal (right) examples generated by the SSM. Knotpoints are indicated by black circles; lines are for the purposes of illustration.

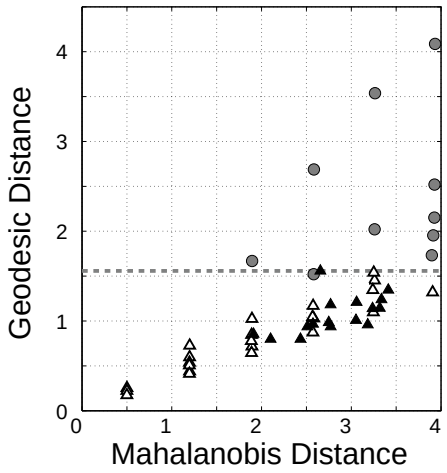

Figure 3: Mahalanobis vs. geodesic distances from the mean shape for *grey circles:* illegal shapes generated by the SSM, *white triangles:* legal shapes generated by the SSM, *black triangles:* the training set.

4.2 Local Analysis

By local analysis, we mean the investigation of a single image warp. For example, we may be able to determine that a given warp is in some sense the least representative of a set of warps; we would then like to be able to say what it is in the related image that is causing this difference. Previous authors have used the Jacobian of the deformation field, which gives the local change of scale; however, what is potentially more interesting are the places where the Jacobian changes. This suggests that instead we should look at curvature measures of the deformation field.

The warp energy itself is an integral over the image of some curvature measure. In the biharmonic case, the warp energy is the approximate Willmore energy, where the Willmore energy is the square of the mean curvature of the deformation field. We could also investigate the Gaussian curvature. All these curvature measures involve second derivatives of the deformation field, which, for certain choices of Green's function (e.g., the biharmonic case), will diverge at the knotpoints.

We can also consider differential-geometric scalar invariants. In n dimensions, we have the source space, with Cartesian coordinates $\{x_\mu\}$ and flat metric $\delta_{\mu\nu}$, that maps to the point $\vec{\Phi}(\vec{x}) = \{X_\alpha(\vec{x})\}$ in the target space. This map induces a second metric on the source space:

$$g_{\mu\nu}(\vec{x}) = \frac{\partial\vec{\Phi}}{\partial x_\mu} \cdot \frac{\partial\vec{\Phi}}{\partial x_\nu} = \left(\partial_\mu\vec{\Phi}\right) \cdot \left(\partial_\nu\vec{\Phi}\right). \tag{13}$$

The Jacobian of the mapping is just the determinant of this induced metric. We might be tempted to calculate the Ricci scalar for this metric; however, since this is just a reparameterisation of the flat metric on the target space, hence the curvature is identically zero.
We can construct the vector-valued tensor:

$$\vec{\Phi}_{\mu\nu} = \partial_\mu\partial_\nu\vec{\Phi}, \tag{14}$$

which can give the approximate Willmore energy (thin-plate bending energy) of the warp:

$$\left(\mathrm{Tr}\vec{\Phi}_{\mu\nu}\right) \cdot \left(\mathrm{Tr}\vec{\Phi}_{\mu\nu}\right). \tag{15}$$

Given that we have two Riemannian metrics on the source space, we can also consider the difference of these two metrics. This is the pseudo-Riemannian metric:

$$h_{\mu\nu} = g_{\mu\nu} - \delta_{\mu\nu}. \tag{16}$$

Muñoz Masqué and Valdés Morales (1994) show that the lowest-order scalar invariant that can be built from a pseudo-Riemannian metric in n dimensions is of order 2 (that is, involves second derivatives). In 2D, there is only one such invariant, the Ricci Scalar, whereas in 3D there are 3 second-order invariants. It is trivial, but tedious, to calculate the Ricci Scalar for the metric $h_{\mu\nu}$.

5 Conclusions

The analysis of the results of non-rigid registrations can be considered as a problem in differential geometry, inasmuch as we are imposing a metric on the diffeomorphism group. We have demonstrated that the approach of generating metrics based on the spline basis gives superior results to the use of ad hoc Euclidean metrics.

We have also shown how ideas from differential geometry can be applied to the local analysis of a single warp.

References

T. Boggio. Sulle funzioni di green d'ordine m. *Rendiconti - Circolo Matematico di Palermo*, 20:97–135, 1905.

F. L. Bookstein. Principal Warps: Thin-Plate Splines and the Decomposition of Deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989.

V. Camion and L. Younes. Geodesic interpolating splines. In M. Figueiredo, J. Zerubia, and A. K. Jain, editors, *Proceedings of EMMCVPR'01*, volume 2134 of *Lecture Notes in Computer Science*, pages 513–527. Springer, 2001.

T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models – their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

J. Duchon. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *Revue Française d'Automatique, Informatique, Recherche Opérationelle (RAIRO) Analyse Numerique*, 10:5–12, 1976.

P. Dupuis, U. Grenander, and M. I. Miller. Variational problems on flows of diffeomorphisms for image matching. *Quarterly of Applied Mathematics*, 56(3):587–600, 1998.

S. C. Joshi and M. M. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Transactions on Image Processing*, 9(8):1357–1370, 2000.

L. LeBriquer and J. Gee. Design of a statistical model of brain shape. In *Proceedings of IPMI'97*, Lecture Notes in Computer Science 1230, pages 477–482, 1997.

S. Marsland and C. J. Twining. Constructing data-driven optimal representations for iterative pairwise non-rigid registration. To be presented at WBIR 2003.

J. Muñoz Masqué and A. Valdés Morales. The number of functionally independant invariants of a pseudo-Riemannian metric. *Journal of Physics A: Mathematical and General*, 27:7843–7855, 1994.

S. M. Pizer, D. S. Fritsch, P. Yushkevich, V. Johnson, and E. Chaney. Segmentation, registration, and measurement of shape variation via image object shape. *IEEE Transactions on Medical Imaging*, 18 (10):851–865, 1999.

D. Rueckert, A. F. Frangi, and J. A. Schnabel. Automatic construction of 3D statistical deformation models using non-rigid registration. In *Proceedings of MICCAI'01*, volume 2208 of *Lecture notes in Computer Science*, pages 77–84, 2001.

C. J. Twining and S. Marsland. Constructing diffeomorphic representations of non-rigid registrations of medical images. In *Proceedings of IPMI 2003*, 2003.

C. J. Twining, S. Marsland, and C. J. Taylor. Measuring geodesic distances on the space of bounded diffeomorphims. In *Proceedings of BMVC'02*, volume 2, pages 847–856. BMVA Press, 2002.