A Non-Euclidean Metric for the Classification of Variations in Medical Images

Carole Twining* and Stephen Marsland*

Imaging Science and Biomedical Engineering, University of Manchester, Manchester M13 9PT, UK

Abstract. The analysis of deformation fields, such as those generated by non-rigid registration algorithms, is central to the quantification of normal and abnormal variation of structures in registered images. The correct choice of representation is an integral part of this analysis. This paper presents methods for constructing multi-dimensional diffeomorphic representations of deformations. We demonstrate that these representations are suitable for the description of medical image-based deformations in 2 and 3 dimensions. We show (using a set of 2D outlines of ventricles) that the non-Euclidean metric inherent in this representation is superior to the usual ad hoc Euclidean metrics in that it enables more accurate classification of legal and illegal variations.

1 Introduction

Non-rigid registration algorithms [1, 2] automatically generate dense (i.e., pixel-to-pixel or voxel-to-voxel) correspondences between pairs and sets of images with the aim of aligning analogous 'structures'. The deformation fields implicit in this correspondence contain information about the variability of structures across the set, and in order to analyse quantitatively this variability, we need to be able to analyse the set of deformation fields. Such analysis must be based (either implicitly or explicitly) on a particular mathematical representation of the deformation field. Previous work on modelling dense 2D and 3D deformation fields has either used the densely-sampled deformation vectors directly (e.g., [3, 4]), or has employed a smooth, continuous representation of these (e.g., [5]). However, neither of these methods guarantees that the deformation field is diffeomorphic.

We contend that the appropriate representation should be continuous and diffeomorphic. Where such a correspondence is not actually physically meaningful (e.g., in the case where additional structures such as tumours appear), this should be indicated by the warp parameters assuming atypical values. When we are considering the correspondence between discrete and bounded objects such as brains, it is also desirable that the warps themselves should be discrete and bounded. This leads us to suggest that a suitable representation is that of the group of continuous diffeomorphisms with some appropriate set of boundary conditions. Such a representation can be constructed using an approach based on Geodesic Interpolating Splines (GIS) [6]. In previous work [6, 7] it has been shown that this approach also allows the construction of a metric on the diffeomorphism group.

In this paper we demonstrate the construction of these diffeomorphic representations using a variety of spline bases. We show that these representations generate warps that are suitable for the task in hand, giving biologically 'plausible' warps in both two and three dimensions, whilst being of a relatively low dimensionality. We further study the significance of the metric (geodesic) distances between warps, and show that using it provides a measure of atypical variation that has greater discriminatory power than naïve measures based on the ad hoc use of a Euclidean metric on the space of warp parameters.

2 The Geodesic Interpolating Spline

We consider a vector-valued spline function $\vec{f}(\vec{x}),\ \ \vec{x} \in \mathbb{R}^n$ that interpolates between data values at a set of knotpoints $\{\vec{x}_i : i = 1 \text{ to } N\}$, where $\vec{f}(\vec{x}_i) = \vec{f}_i$, that can be expressed as the minimiser of a functional Lagrangian of the form:

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

where L is some scalar differential operator. The first term in the Lagrangian is the smoothing 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 in the form:

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

where the affine function g is a solution of $L\vec{g}(\vec{x}) = 0$ and the Green's function G is a solution of: $\left(L^\dagger L\right) G(\vec{x}, \vec{y}) \propto \delta(\vec{x} - \vec{y})$, with $L^\dagger$ is the Lagrange dual of L. For more details, see [8]. The choice of Green's function depends on the boundary conditions and smoothness appropriate to the problem considered. Suggestion of different possible Green's functions are given in [8], here we focus on the Clamped-Plate Spline (CPS), which has the boundary conditions that it is identically zero on and outside the unit ball [9]. We contend that such boundary conditions are the appropriate choice for images of discrete objects such as brains; other types of images may require different boundary conditions. The biharmonic CPS ($L^\dagger L = (\nabla^2)^2$) in 2 dimensions G_2^2, and the triharmonic ($L^\dagger L = (\nabla^2)^3$) CPS in 3 dimensions G_3^3, have Green's functions ([9, 10]):

*Joint first authors. {carole.twining,stephen.marsland}@man.ac.uk

$G_2^2(\vec{x},\vec{y}) = \|\vec{x}-\vec{y}\|^2\left(A^2 - 1 - \log A^2\right), G_3^3(\vec{x},\vec{y}) = \|\vec{x}-\vec{y}\|\left(A + \frac{1}{A} - 2\right), A(\vec{x},\vec{y}) = \frac{\sqrt{\vec{x}^2\vec{y}^2 - 2\vec{x}\cdot\vec{y}+1}}{\|\vec{x}-\vec{y}\|}$.
The CPS is only guaranteed to be diffeomorphic for infinitesimal deformations. The standard approach to constructing larger deformations is to build them up as an infinite sequence of infinitesimal deformations [6, 11] by introducing a flow time t, so that the knotpoints follow paths $\{\vec{x}_i(t); t \in [0,1]\}$ with the associated energy:

$$E\left[\vec{\alpha}_k(t)\right] = \int_0^1 dt\, G\left(\vec{x}_i(t), \vec{x}_j(t)\right)\left(\vec{\alpha}_i(t)\cdot\vec{\alpha}_j(t)\right). \tag{3}$$

We no longer have an exact solution, since the knotpoint paths are only constrained at their end-points, so that we have to numerically optimise equation (3) over the knotpoint paths between their end-points. For more details, see [7]. It was shown in [6] that the optimised energy is the square of a *geodesic distance function* d on the group of diffeomorphisms, so that $E_{\text{opt}}(\omega) = d^2(e,\omega)$, where e is the identity element of the group. This metric gives us a principled way of defining warps that interpolate between any two given warps [10]; the optimal flowpath in the group of diffeomorphisms gives a geodesic on the space of warps, and the geodesic distance allows us to calculate a warp on this geodesic halfway between the two initial warps.

3 Representing Diffeomorphisms in Two and Three Dimensions

When considering warps of 2D biological images, it is obviously important that the generated warps are not only diffeomorphic, but also biologically plausible. To investigate this, we considered a set of 2D MR axial slices of brains, where the slices chosen 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 knots (see Fig. 1). Given a pair of images, the knotpoint positions on the images gave us the initial and final positions for our knotpoint paths. We then calculated the geodesic interpolating spline warp corresponding to these positions using the 2D clamped-plate spline as Green's function.

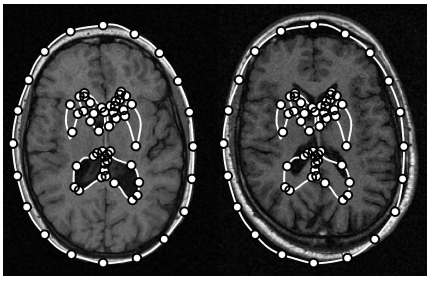

Figure 1. *Left:* Annotation (white line) and knotpoints (white circles) on the original brain slice. *Right:* The same knots positioned on another brain slice.

We did not affinely align the knots before calculating the warp; hence the algorithm had to deal with a non-trivial pseudo-affine part. Affine alignment *could* have been performed first, but we did not in order to make the problem *harder*. Example results are shown in Fig. 2. The warped images are not resampled – the images are instead plotted as coloured surfaces, so that the size and position of each warped pixel is retained. It can be seen that the warps are indeed diffeomorphic, and appear to be very smooth – each of the brain slices still looks biologically plausible, despite the relatively low dimensionality of the representation used – structures other than the labelled ones have been brought into approximate alignment. This suggests that a dense correspondence (for instance, one given by a non-rigid registration using maximisation of mutual information) could also be represented by these warps without an inordinate increase in the dimensionality of the representation.

We now show that the GIS can also generate biologically warps in 3D, and that, given a warp, we can choose the knotpoints appropriately using a set of segmented hippocampi, each of which consists of a triangulated surface with 268 vertices; examples are shown in Fig. 3. The vertices have been manipulated to give the optimisal correspondence [12]. Pairs of hippocampi were chosen at random, and the 2 shapes aligned using generalised Procrustes analysis. We used the triharmonic clamped-plate spline G_3^3 as our GIS basis [8]. The required warp between source and target was calculated iteratively – the warp was optimised for a given set of knotpoints, then new knotpoints added and the warp recalculated. New knotpoints were selected from the vertices using a greedy algorithm: the discrepancy between the vertices of the warped source and the target were calculated and new knotpoints selected from those vertices that have the largest discrepancies.

Fig. 4 shows the distribution of the discrepancies between the aligned source and target, and the final warped source and target, for a set of 70 knotpoints. It can be seen that the distribution of discrepancies as a whole has been shifted towards smaller values. In Fig. 5, we show the maximum, median and mean square discrepancies for non-knot points only as a function of the number of knots for 4 random pairs of hippocampi. The nature of our greedy algorithm for selecting knotpoints means that the maximum discrepancy is not guaranteed to decrease monotonically. However, all three graphs show that the algorithm quickly reaches to a reasonable representation of the required warp, for a number of knotpoints that is approximately 25% of the number of vertices.

4 Using the Geodesic Distance to Classify Variations

We now consider the role of the geodesic distance in classifying legal and illegal variations in real biological data. We take as our dataset the annotated outlines of the anterior lateral ventricles, as used in section 3 in the axial brain slices. Each example consists of 40 knotpoints (see Fig. 6). The set of training examples was Procrustes aligned

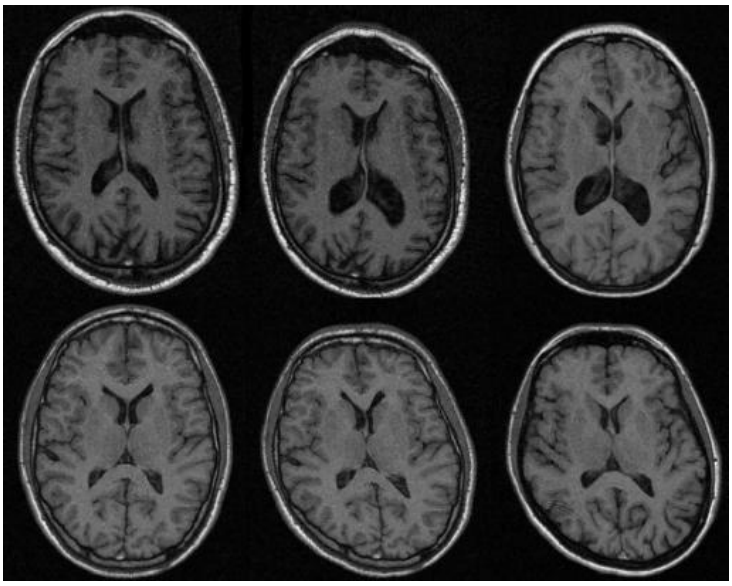

Figure 2. Two examples of warp interpolation using the clamped-plate spline. Pixel intensity is unchanged, but note that the image structures are approximately aligned. *Left:* Source image, *Centre:* Warped image, *Right:* Target image. Source and target images are *undistorted* images from 4 normal subjects.

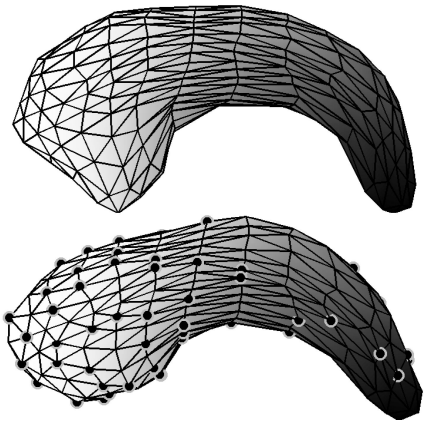

Figure 3. Target (top) and source (bottom) hippocampi with knotpoints (black circles). The correspondence between the shapes is indicated by the shading.

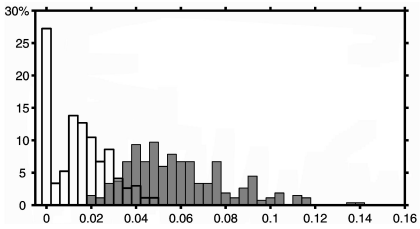

Figure 4. Distribution of point discrepancy between source and target (grey bars), and warped source and target (white bars).

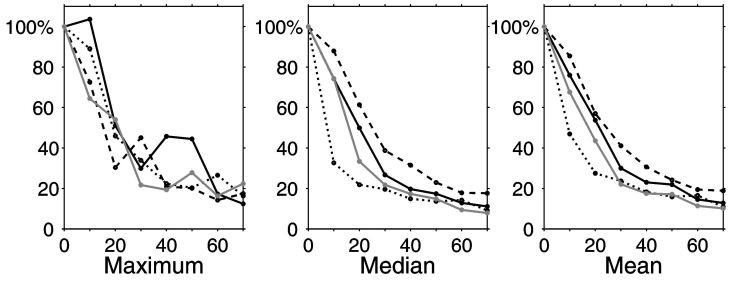

Figure 5. The maximum, median and mean square discrepancies, for non-knot points only, as a function of the number of knots. Data is shown from 4 randomly selected pairs of hippocampi.

and then scaled to fit inside the unit circle. A linear Statistical Shape Model (SSM) was built from this training set in the usual way. 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 (see Fig. 6). The training set of shapes are, by definition, legal.

We then calculated the GIS warps, using the biharmonic CPS basis, between the classified set of shapes and the mean shape from the model. The geodesic distance from the mean is compared with the Mahalanobis distance from the mean in Fig. 7. It is immediately obvious that we cannot separate the legal and illegal shapes by using the Mahalanobis distance from the mean. 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 again demonstrates the superiority of the GIS metric over an ad hoc metric. Note that the correspondences used in this example are a subset of the correspondences that we would expect to be generated by a successful non-rigid registration of the images. Increasing the density of points on the training shapes would have left the result for the Mahalanobis distance essentially unchanged. However, the result for the GIS warp would have improved, giving a greater separation between the two sets of shapes. This is because, in the limit where the lines become infinitely densely sampled, it is actually impossible to construct a diffeomorphism for which the lines cross, which would mean that the geodesic distance for the illegal shapes would approach infinity as the sampling density increased. We can now extend this argument to the case of modelling the deformation fields for a non-rigid registration; a linear model of such deformation fields would suffer the same problem as the linear SSM, where now the overlapping structures would correspond to a folding of the warp. The GIS cannot, by definition, generate such a folding since it is guaranteed to be diffeomorphic.

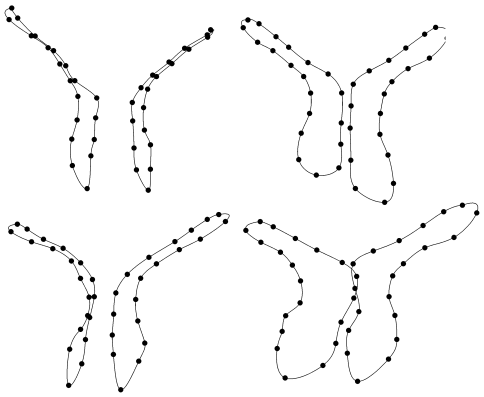

Figure 6. *Top:* Examples from the training set. *Bottom:* Legal (left) and illegal (right) examples generated by the SSM. Knotpoints are indicated by black circles; lines are for the purposes of illustration only.

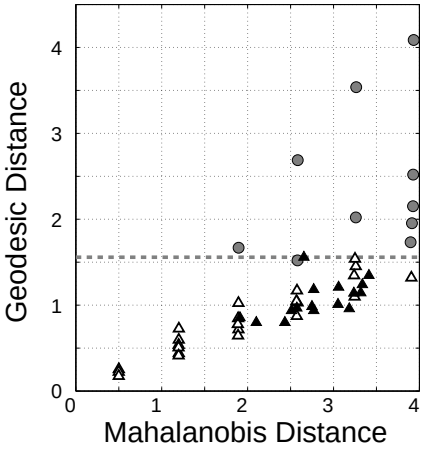

Figure 7. Mahalanobis vs. geodesic distances from the mean shape for *Grey circles:* illegal shapes generated by the SSM, *White triangles:* legal shapes generated by the SSM, *Black triangles:* the training set.

5 Conclusions

This paper has introduced a principled diffeomorphic representation of deformation fields with an inherent non-Euclidean metric; the spline basis of this representation is defined by the choice of Green's function and boundary conditions, which can be altered to suit the particular task in hand. We have demonstrated that this representation method can accurately represent real biological variations in both two and three dimensions. Conventional linear modelling strategies impose a Euclidean metric on the space of parameters (in our case, the knotpoint positions). The Mahalanobis distance that we have used for comparisons in this paper is derived from such a metric. The example in section 4 clearly shows the superiority of the non-Euclidean metric in quantifying variation.

Acknowledgements

Our thanks to G. Gerig and M. Steiner for the hippocampus dataset, and to Rh. Davies for providing us with his optimised correspondences for this dataset.

References

1. C. Chefd'Hotel, G. Hermosillo & O. Faugeras. "A variational approach to multi-modal image matching." In *Proceedings of IEEE Workshop on Variational and Level Set Methods (VLSM'01)*, pp. 21 – 28. 2001.
2. D. Rueckert, L. I. Sonoda, C. Hayes et al. "Non-rigid registration using free-form deformations: Application to breast MR images." *IEEE Transactions on Medical Imaging* **18(8)**, pp. 712–721, 1999.
3. A. Guimond, J. Meunier & J.-P. Thirion. "Average brain models: A convergence study." Technical Report RR-3731, INRIA, Sophia Antipolis, 1999.
4. L. LeBriquer & J. Gee. "Design of a statistical model of brain shape." In *Proceedings of IPMI'97*, volume 1230 of *Lecture Notes in Computer Science*, pp. 477–482. Springer, 1997.
5. D. Rueckert, A. F. Frangi & 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*, pp. 77–84. 2001.
6. V. Camion & L. Younes. "Geodesic interpolating splines." In M. Figueiredo, J. Zerubia & A. K. Jain (editors), *Proceedings of EMMCVPR'01*, volume 2134 of *Lecture Notes in Computer Science*, pp. 513–527. Springer, 2001.
7. C. J. Twining, S. Marsland & C. J. Taylor. "Measuring geodesic distances on the space of bounded diffeomorphims." In *Proceedings of BMVC'02*, volume 2, pp. 847–856. BMVA Press, 2002.
8. C. J. Twining & S. Marsland. "Constructing diffeomorphic representations of non-rigid registrations of medical images." In *Proceedings of IPMI, to appear*. 2003.
9. T. Boggio. "Sulle funzioni di green d'ordine m." *Rendiconti - Circolo Matematico di Palermo* **20**, pp. 97–135, 1905.
10. S. Marsland & C. J. Twining. "Clamped-plate splines and the optimal flow of bounded diffeomorphisms." In *Statistics of Large Datasets, Proceedings of Leeds Annual Statistical Research Workshop*, pp. 91–95. 2002.
11. P. Dupuis, U. Grenander & M. I. Miller. "Variational problems on flows of diffeomorphisms for image matching." *Quarterly of Applied Mathematics* **56(3)**, pp. 587–600, 1998.
12. R. Davies, C. J. Twining, T. F. Cootes et al. "3D statistical shape models using direct optimisation of description length." In *Proceedings of ECCV*, volume 2352 of *Lecture Notes in Computer Science*, pp. 3–20. Springer, 2002.