

Clamped-Plate Splines and the Optimal Flow of Bounded Diffeomorphisms

Stephen Marsland and Carole Twining

Division of Imaging Science and Biomedical Engineering
University of Manchester, Oxford Road, Manchester M13 9PT

Email: stephen.marsland@man.ac.uk, carole.twining@man.ac.uk

1 Introduction

Non-rigid registration of biomedical images has been used as a way of establishing a meaningful dense correspondence between images of the same physical object. For example, it can be used to establish a dense correspondence between images of the same object taken using different imaging modalities (e.g., MRI and X-ray), or between a new example and a previously constructed annotated anatomical atlas.

In general, a dense set of points on a sample image are ‘warped’ by a (non-linear) function, so that the reconstructed object in one image has the ‘same’ appearance as the object in the target image. Given a set of images that have been brought into mutual registration (by whatever method), the set of warps implicitly encodes information about the variability of the structures present in the images. If the warping functions are constrained to be smooth whilst also not tearing or folding the image, then the functions are bijective, invertible and differentiable to some order. Furthermore, if the objects considered are discrete and bounded, it seems that the appropriate set of warp functions to consider will belong to the group of diffeomorphisms with some non-trivial boundary conditions.

Finally, in order to be able to statistically analyse the variability and similarity of the imaged objects using the information encoded in the warp fields, some method of measuring a well-defined ‘distance’ on the space of warps is required. Given some parameterized representation of the warps, it should be possible to work directly in this space of parameters. However, it is clear that imposing an ad hoc Euclidean distance on this parameterised space cannot be considered meaningful in any geometric or group-theoretic sense.

This paper aims to develop general tools for building statistical models on the space of warps of pixellated/voxellated images. In previous work (e.g., Camion and Younes (2001)) the idea of a metric distance between warps was used in the context of **inexact** landmark matching, where the closeness of fit of the set of landmarks was balanced against the metric distance of the warp from the identity. For our application, we consider that a previously generated warp of a pixellated image is completely defined by the transformed coordinates of all the image pixels; these then become our ‘landmarks’. We are hence only interested in the exact matching case, since this corresponds to the exact description of the warp.

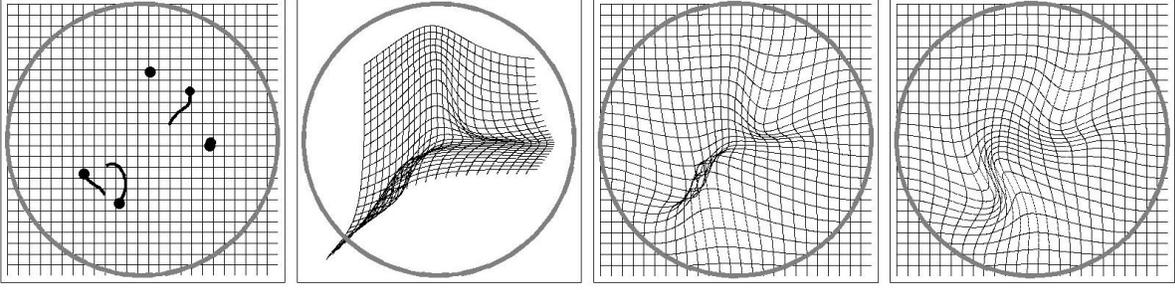


Figure 1: From left to right: The original image, the paths of five (randomly chosen) control points and the unit circle (grey), the thin-plate spline interpolant of the deformation field, the clamped-plate spline interpolant and the geodesic interpolating spline.

2 Metrics, Splines and Diffeomorphisms

2.1 A Suitable Metric

To be suitable for our purposes the metric distance between two warps should be independent of the choice of reference warp. That is, the metric on the space of diffeomorphisms must be invariant under the action of the group $\text{Diff}(M)$, i.e.,

$$d(\theta, \phi) = d(\psi \circ \theta, \psi \circ \phi) \quad \forall \theta, \phi, \psi \in \text{Diff}(M),$$

where $\psi \circ \theta$ is the group multiplication. Camion and Younes (2001) show how to construct metrics that have this property. They describe an energy minimization algorithm that enables them to compute a geodesic flow of diffeomorphisms on deformable data. Their implementation is based on interpolating splines such as ‘thin-plate splines’ (Bookstein, 1989) and Gaussian splines that do not have any specific boundary conditions imposed.

We extend their method by introducing a spline with both Dirichlet and von Neumann boundary conditions, which we will call the ‘clamped-plate spline’. This seems an appropriate choice for the set of warps in the case of discrete objects; the mapping of the boundary between different examples corresponds to performing an initial **rigid** registration of the images.

2.2 The Clamped-Plate Spline

Consider a function $f(\mathbf{r})$, $\mathbf{r} \in \mathbb{R}^n$, with n_c constraints $\{f(\mathbf{r}_i) = f_i, i = 1, \dots, n_c\}$. The clamped-plate spline interpolant for this function has the form $f(\mathbf{r}) = \sum_i \alpha_i G_{(n)}(\mathbf{r}, \mathbf{r}_i)$, where

$G_{(n)}(\mathbf{r}, \mathbf{R})$ is the Greens function of the biharmonic equation in n -dimensions with Dirichlet and von Neumann boundary conditions (Boggio, 1905):

$$G_{(n)}(\mathbf{r}, \mathbf{R}) = |\mathbf{r} - \mathbf{R}|^{4-n} \int_1^{A(\mathbf{r}, \mathbf{R})} \frac{(v^2 - 1)}{v^{n-1}} dv, \quad A(\mathbf{r}, \mathbf{R}) = (|\mathbf{r} - \mathbf{R}|)^{-1} \sqrt{|\mathbf{r}|^2 |\mathbf{R}|^2 - 2\mathbf{r} \cdot \mathbf{R} + 1}.$$

This interpolant minimizes the same approximate form of the Willmore (1992) energy (the bending-energy) as the thin-plate spline (Bookstein, 1989), but has the important property that the function and its first derivative are zero on the boundary of the unit ball \mathbb{D}^n .

2.3 Geodesic Interpolating Splines (GIS)

Consider a set of control points in \mathbb{D}^n with initial and final positions $\{\mathbf{q}_i(0), \mathbf{q}_i(1)\}$. Under a flow of diffeomorphisms these points trace out paths $\{\mathbf{q}_i(t)\}$ that are linked to the associated velocity field by

$$\frac{d\mathbf{q}_i(t)}{dt} \equiv \mathbf{v}(t, \mathbf{q}_i(t)).$$

The energy associated with this diffeomorphic flow is given by the integrated bending energy for the velocity fields, which are expanded in terms of the clamped-spline interpolant. Calculating a **geodesic** flow of diffeomorphisms with constraints hence corresponds to analytically optimising this bending energy term with appropriate constraints:

$$E = \int_0^1 dt \int_{\mathbb{R}^n} d\mathbf{q} |L\mathbf{v}(t, \mathbf{q}(t))|^2 + \lambda \sum_{i=1}^{n_c} \int_0^1 dt \left| \mathbf{v}(t, \mathbf{q}_i(t)) - \frac{d\mathbf{q}_i(t)}{dt} \right|^2, \quad L = \nabla^2$$

where the sum is over the number of control points n_c and λ is a Lagrange multiplier. Camion and Younes (2001) give such an algorithm for the case of inexact landmark matching, while for our case of exact matching we explicitly impose the constraint

$$\mathbf{v}(t, \mathbf{q}_j(t)) = \frac{d\mathbf{q}_j(t)}{dt} = \sum_{i=1}^{n_c} \alpha_i(t) G(\mathbf{q}_i(t), \mathbf{q}_j(t)),$$

and so the metric distance in the space of diffeomorphic flows is

$$d[\{\mathbf{q}_i(t)\}] = \sum_{i=1}^{n_c} \int_0^1 dt \alpha_i(t) \cdot \mathbf{v}(t, \mathbf{q}_i(t)).$$

Hence, for an optimal set of control point paths $\{\mathbf{q}_i(t)\}$, which minimise the above metric distance, the optimal velocity field is fully determined and the path of any arbitrary point can be computed by integration. This then defines the action of a GIS warp parameterised by the initial and final positions of the control points. We have implemented an algorithm that computes these optimal control point paths for 1, 2 and 3 dimensions. The outputs of the algorithm have been compared with exact theoretical results for each case, and the algorithm has been shown to quickly converge to the exact result and to have the correct time symmetry properties.

3 Generating Arbitrary Diffeomorphic Warps and Interpolating Between Them

Let $\omega \equiv \omega(\{\mathbf{q}_i(0)\}, \{\mathbf{q}_i(1)\}) \in \text{Diff}(\mathbb{D}^n)$ denote a geodesic interpolating spline (GIS) diffeomorphism. Given any arbitrary bounded diffeomorphism g that acts on some dense pixel set S_e , a GIS approximant $\omega(S_e) \approx g(S_e)$ can be iteratively constructed in the following manner (with $\omega = e$, the group identity element and $n_c = 0$ being the initial approximant):

- Find the pixel position $\mathbf{y} \in S_e$ where the difference $|\omega(\mathbf{y}) - g(\mathbf{y})|$ is maximal
- Add this point to the existing set of control points of ω , with endpoints \mathbf{y} and $g(\mathbf{y})$
- Update the warp ω
- Iterate until convergence

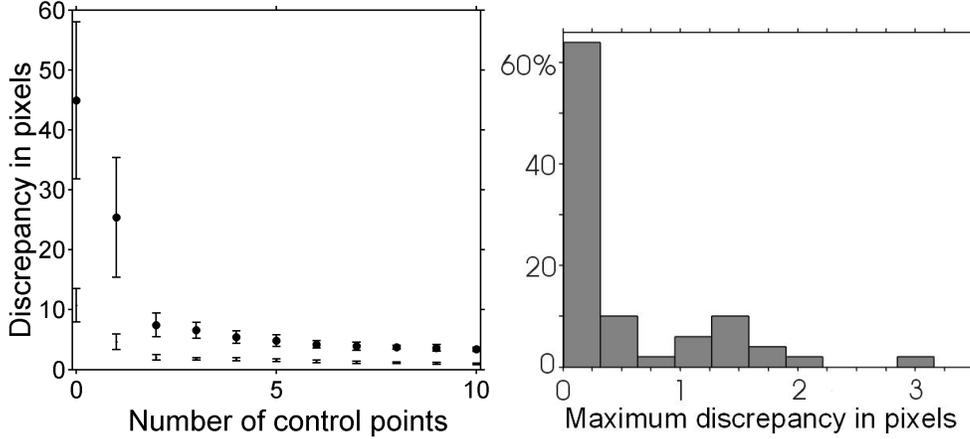


Figure 2: *Left:* The maximum (\bullet) and mean (\cdot) pixel discrepancies between a warp and its approximant as a function of n_c . *Right:* The distribution of maximum discrepancies in pixels for the interpolated warp.

The left of figure 2 shows the maximum and mean pixel discrepancies between a diffeomorphic warp of \mathbb{D}^2 and its approximant as the number of control points increases. The exact diffeomorphic warps used for testing were created using the cumulative distribution of a wrapped Cauchy function (Mardia, 1972), with several such warps with random parameters being concatenated. The graph shows that the approximant converges rapidly.

As any arbitrary bounded diffeomorphism can be approximated to any required degree of accuracy by a GIS warp (for a sufficiently large number of control points n_c), for any set of warps $\{g_a : a = 1, \dots, n\}$ we can construct the equivalent set of GIS warps $\{\mu_a\} \approx \{g_a\}$, $\mu_a = \omega(\{\mathbf{q}_i(0)\}, \{\mathbf{q}_i^a(1)\})$. The information about the distribution of the set of warps in the space of diffeomorphisms is encoded by the set of geodesic distances between all pairs of warps:

$$d(\mu_a, \mu_b) \equiv d(e, \mu_a \circ \mu_b^{-1}) \equiv d(e, \mu_b \circ \mu_a^{-1}),$$

where $\mu_a^{-1} = \omega(\{\mathbf{q}_i^a(1)\}, \{\mathbf{q}_i(0)\})$ is the inverse warp. We therefore need to be able to construct warps that interpolate between μ_a and μ_b , that is, warps of the form $\mu_b \circ \mu_a^{-1}$. These warps can be approximated by the warp $\nu_{ab} = \omega(\{\mathbf{q}_i^a(1)\}, \{\mathbf{q}_i^b(1)\})$, which is exactly equivalent to $\mu_b \circ \mu_a^{-1}$ at the control points. The accuracy of the approximant between the control points was computed by comparing $\nu_{ab}(\mu_a(S_e))$ and $\mu_b(S_e)$, since $(\mu_b \circ \mu_a^{-1})(\mu_a(S_e)) \equiv \mu_b(S_e)$. The maximum discrepancies in pixels between each interpolated warp pair from a set of 50 test warps is shown on the right of figure 2. In 80% of cases the maximum discrepancy is less than one pixel.

4 Applications of the Geodesic Distance

We have applied the above techniques to the analysis of diffeomorphic warps derived from two-dimensional images of real biological objects. The behaviour of the geodesic distances across the set of warps was compared to the Mahalanobis distances calculated from a linear model based on the space of warp parameters. It was seen that the geodesic distance makes it easier to distinguish between biologically plausible and implausible object deformations (see figure 3).

Furthermore, we have considered the case of classification of variation in biological objects; here the geodesic distances were used as the input to a Support Vector Machine classifier.

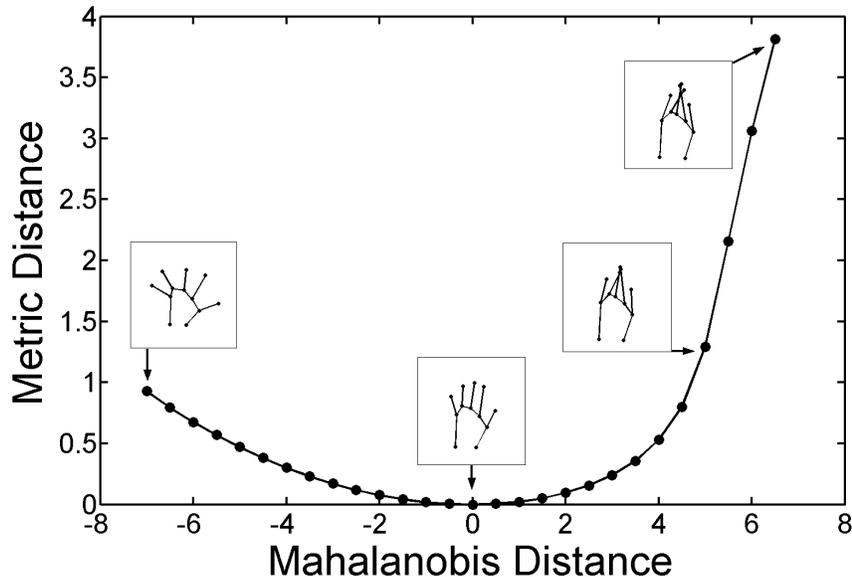


Figure 3: Mahalanobis distance versus geodesic distance for warps based on the variation of a set of hand outlines. Note that the geodesic distance penalizes those warps for which the fingers cross, a variation that was not present in the training set.

5 Conclusions and Future Work

This paper has described a well-defined metric on the space of diffeomorphic warps. We have introduced a new spline, the clamped-plate spline, which enables us to construct arbitrary bounded diffeomorphisms on the unit ball \mathbb{D}^n . Using the clamped-plate spline, an algorithm has been developed that computes the geodesic flow of diffeomorphisms; the algorithm has been verified by comparison with exact theoretical results in 1, 2 and 3 dimensions.

We have shown that we can approximate an arbitrary bounded diffeomorphic warp to any required degree of accuracy and interpolate between any two arbitrary warps, which enables us to calculate the geodesic distances between all members of a set of bounded diffeomorphic warps. Preliminary results have been presented that show the application of the method to the analysis of sets of warps derived from real data.

Future work will consider applying the method to larger datasets, including those in three-dimensions, and explicit model building on the space of diffeomorphic warps rather than warp parameters.

References

- T. Boggio. Sulle funzioni di green d'ordine m . *Circolo Matematico di Palermo*, 20:97–135, 1905.
- Fred 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.
- Vincent Camion and Laurent Younes. Geodesic interpolating splines. In *Proceedings of Energy Minimization in Computer Vision and Pattern Recognition (EMMCVPR)*, volume 2134 of *Lecture Notes in Computer Science*, pages 513 – 527. Springer-Verlag, 2001.
- K. V. Mardia. *Statistics of Directional Data*. Academic Press, 1972.
- T. J. Willmore. A survey on Willmore immersions. In *Geometry and Topology of Submanifolds, IV*, pages 11–16. World Scientific Publishers, 1992.