Groupwise Non-Rigid Registration using Polyharmonic Clamped-Plate Splines

Stephen Marsland[1], Carole J. Twining[1] and Chris J. Taylor

Imaging Science and Biomedical Engineering,
University of Manchester, Manchester M13 9PT, U.K.

Abstract. This paper introduces a novel groupwise data-driven algorithm for non-rigid registration. The motivation behind the algorithm is to enable the analysis of groups of registered images; to this end, the algorithm automatically constructs a low-dimensional, common representation of the warp fields. We demonstrate the algorithm on an example set of 2D medical images, and show that we can obtain good registration across the set, with automatic detection and correction of misaligned examples, whilst still maintaining a low-dimensional representation.

1 Introduction

The registration of sets of medical images, i.e., finding dense correspondences between images with the aim of aligning analogous structures, is an important problem in medical image analysis [3, 5, 8, 15]. Tasks that are suitable for image registration include the correction for patient motion either during a single scan sequence (e.g., [15]), or between scans of the same patient taken at different times, such as in the monitoring of disease progression. Other common examples are the registration of multi-modal (e.g., MR, PET and CT) images for either a single patient or multiple patients [16], the registration of time-series of a moving object such as the heart [4], and registration onto an anatomical atlas for automated segmentation [8].

The majority of registration algorithms proposed in the literature are based on aligning pairs of images. However, 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 – the importance of representation in modelling has been demonstrated by the fact that explicitly optimizing the representation can lead to appreciable improvements in model performance [6]. It is this requirement of a common mathematical representation that underlies the method of groupwise non-rigid registration that is introduced in this paper.

We contend that groupwise registration should be considered as a separate problem in image registration, not just as a series of pairwise registrations. While pairwise algorithms can provide accurate registration between two images, a series of pairwise registrations will not necessarily have a common underlying

[1] Joint first authors. Contact: {stephen.marsland, carole.twining}@man.ac.uk

representation; conversely, where there is a common representation it may be too general (and too high-dimensional) for direct use in analysis. These problems can be seen in previous work on modelling based on sets of pairwise registrations, which have either used the densely-sampled deformation vectors directly [9, 10], or employed a smooth, continuous representation of them [14].

In this paper we describe a data-driven algorithm for registering, iteratively, a group of images in 2D or 3D. A common representation of the deformation fields, based on interpolating splines, is introduced; the set of knotpoints that define the splines are refined as the algorithm iterates in order to improve all of the registrations in the group. In addition, the frame of reference (i.e., the choice of reference image) can also change in order to increase the overall accuracy of the registrations. We demonstrate the action of the algorithm on a set of 2D MR images of axial slices of human brains.

2 Spline-Based Representations and Polyharmonic Clamped-Plate Splines

The non-rigid registration algorithm described in this paper is based on the polyharmonic (order m) interpolating clamped-plate spline (CPS) in n dimensions [18], which is used to interpolate the warped image between knotpoints. A related algorithm has been applied to pairwise registration in [11]. While other algorithms have used spline-based techniques for representing deformation fields, (e.g., the Free-Form Deformations used by Rueckert et al. [15], based on B-splines) such splines are in essence smoothing splines – while the splines can be iteratively refined in a data-driven way by locally increasing the resolution of the mesh of spline knotpoints [12], this requires a computationally expensive algorithm such as the Oslo algorithm [7] and, as observed by Rohlfing and Maurer [12] "it is difficult to preserve consistency". Interpolating splines do not have this drawback – new knotpoints that improve the registration can be introduced without significant additional computational cost, and without affecting any of the other knotpoints.

The polyharmonic clamped-plate spline is represented by the initial and final positions of a set of n_k knots with initial positions $Q_0 = \{\boldsymbol{q}_{0\beta} \in \mathbb{R}^n, \beta = 1, \ldots, n_k\}$ and final positions $Q_1 = \{\boldsymbol{q}_{1\beta}\}$. The associated warp is denoted by $f(\cdot) = \omega(Q_0, Q_1)$, with the action on points $\boldsymbol{x} \in \mathbb{R}^n$ being:

$$f(\boldsymbol{x}) = \boldsymbol{x} + \sum_{\beta=1}^{n_k} \boldsymbol{\alpha}_\beta G_n^m(\boldsymbol{x}, \boldsymbol{q}_{0\beta}). \tag{1}$$

The vector-valued coefficients $\{\boldsymbol{\alpha}_\beta\}$ are found by solving the exact matching conditions for the knotpoints $f(\boldsymbol{q}_{0\beta}) = \boldsymbol{q}_{1\beta}$, and the functions G_n^m are the m^{th} order polyharmonic clamped-plate Green's functions in n dimensions [1]. The boundary conditions on these Green's functions, and hence on the warps, is that the deformation is identically zero on and outside the unit ball in n dimensions, meaning that any part of the image outside the boundary is not warped. In 2 dimensions the biharmonic ($m = 2$) CPS (used in the experiments reported in section 4) has Green's function:

$$G_2^2(\boldsymbol{x},\boldsymbol{y}) = |\boldsymbol{x}-\boldsymbol{y}|^2\left(\frac{1}{2}(A^2-1)-\log A\right),\ \ A(\boldsymbol{x},\boldsymbol{y}) = \frac{\sqrt{|\boldsymbol{x}|^2|\boldsymbol{y}|^2-2\boldsymbol{x}\cdot\boldsymbol{y}+1}}{|\boldsymbol{x}-\boldsymbol{y}|}\,. \tag{2}$$

The clamped-plate bending energy of these warps is given by:

$$E[\omega] = \sum_{\beta,\mu}(\boldsymbol{\alpha}_\beta\cdot\boldsymbol{\alpha}_\mu)\ G_n^m(\boldsymbol{q}_{0\beta},\boldsymbol{q}_{0\mu}), \tag{3}$$

with $d[\omega]=\sqrt{E[\omega]}$ being defined as the associated distance.

3 The Iterative Non-Rigid Registration Algorithm

This section introduces the registration algorithm as a general framework. Various computational considerations are introduced, and the results of using the algorithm, together with further implementation issues, are given in section 4.

3.1 Notation and the Algorithm

We consider a set of images $\{I_i : i = 1,\ldots,N\}$, represented as scalar-valued functions defined on some dense point set X_0. Initially, these points will be a regular grid of points (the set of pixel or voxel centres), which we will denote by X_0. By interpolation we can extend such an image to be continuously defined within the hull of X_0; we can also extrapolate, defining values of I outside this hull as either the background value or the median value of $I(X_0)$, as appropriate.

Images can be warped by applying a warping function $h : X_0 \to X_h = h(X_0)$, to the grid, where the warp could be a general affine warp g, a CPS warp f, or some composition of them, $h = f\circ g$. This warped grid is then used to generate the warped image using the pullback map h^*:

$$h^* : I \to I^{(h)},\ \text{ where } I^{(h)}(X_0) \equiv I(h(X_0)), \Rightarrow I^{(h)}(h^{-1}(X_0)) \equiv I(X_0), \tag{4}$$

which requires interpolation of the function I from the points X_0 to $h(X_0)$.

The similarity between two images I and I' is quantified by the evaluation of some objective function $\boldsymbol{\Phi}\left(I(X_0), I'(X_0)\right)$. One member of the image set I_i is selected as the reference I_s and the remainder are aligned with that image using affine and non-rigid registration, with warps being selected so that:

$$g_i = \arg\min_g \boldsymbol{\Phi}\left(I_s, I_i^{(g)}\right),\ f_i = \arg\min_f \left[\boldsymbol{\Phi}\left(I_s, I_i^{(f\circ g_i)}\right) + \lambda E[f]\right], \tag{5}$$

where $\lambda E[f]$ is a regularisation term proportional to the energy of the warp.

Non-rigid CPS warps are defined by a set of knotpoints, as shown in section 2. The positions of these knotpoints, Q_{orig}, are first defined on the reference image, and then affinely warped, $Q_0^{(i)} = g_i(Q_{\text{orig}})$. The non-rigid warp generated by these knots is $f(q) = \omega(Q_0^{(i)}, q)$, where q are the final knotpoint positions, and the optimised final knotpoints on the free image I_i are then given by:

$$Q_1^{(i)} = \arg\min_q \left[\boldsymbol{\Phi}\left(I_s(X_0), I_i^{(f(q)\circ g_i)}(X_0)\right) + \lambda E[f(q)]\right]. \tag{6}$$

For a set of free images, we define the warp distance between pairs by first affinely registering each pair:

$$g_{ij} = \arg\min_g \boldsymbol{\Phi}\left(I_i(X_0), I_j^{(g)}(X_0)\right), \tag{7}$$

with $g_{si} = g_i$, and then computing the warp distances:

$$d_{ij} = d\left[\omega\left(g_{ij}(Q_1^{(i)}), Q_1^{(j)}\right)\right]. \tag{8}$$

The structure of our iterative algorithm can be summarised as follows, with some features described in more detail afterwards:

The Algorithm

- **Initialisation**
  - **Scale the common pixel/voxel grid X_0 so that it lies within the unit ball**
  - **Select an initial reference image I_S from the set $\{I_i, i = 1, \ldots N\}$**
  - **Pick a set of initial knotpoint positions $Q_{\textbf{orig}}$ on this reference image**
  - **Affinely register every pair of images to obtain $\{g_{ij}\}$**
- **★ Main Loop**
  - **Find the $\{g_i\}$ that affinely align each free image with the reference, and compute the transformed positions of the knotpoints $Q_0^{(i)} = g_i(Q_{\textbf{orig}})$**
  - **Transfer the knotpoints onto the free images to obtain the initial estimate of the final knotpoint positions $Q_1^{(i)} = Q_0^{(i)}$**
  - **Perform pairwise non-rigid registration of each free image to the reference by optimising over the final knotpoint positions $Q_1^{(i)}$**
  - **Calculate the warp distances d_{ij} between every pair of images**
  - **Use this set of distances to select a candidate new reference image $I_{S'}$**
  - **If this differs from the current reference:**
    - **Refine the knotpoint positions $Q_1^{(S')}$ on the candidate reference $I_{S'}$**
    - **Compare the objective function scores for alignment to the original and refined candidate references, accept the new reference if scores improve**
  - **Transfer the knotpoint positions to the frame of the updated reference image**
  - **Use the groupwise discrepancy between the free images and the reference (calculated in the frame of the updated reference) to select new knotpoints**
- **Iterate from ★ until convergence**

Affine Alignment There are two sets of affine alignments computed in the algorithm. Firstly, the free images are each registered onto the chosen reference image. This is the standard pairwise alignment. Secondly, during the first pass through the loop, the alignments between every pair of images are computed to enable the warp distances to be computed and the new reference to be selected. While this is a computationally expensive algorithm ($\mathcal{O}(N^2)$), it only has to be performed once for a given set of images – no matter how many times the algorithm is run – and the optimisation can be initialised using the affine warp given by Procrustes alignment of the knotpoint positions.

Non-Rigid Registration At each pass through the main loop, each of the free images are non-rigidly registered onto the chosen reference. Rather than optimising over all knotpoints simultaneously, we instead optimise over each in turn, and repeat this several times, randomising the order at each repetition. This speeds up the optimisation considerably, without loss of accuracy. Providing that the reference image does not change, a full optimisation is only required for the new knotpoints added during that iteration – the registration is already optimised with respect to the knotpoints defined earlier. Further details about the non-rigid registration between pairs of images, including examples of warps,

can be found in [11]. When the reference image changes, the positions of the knotpoints on that reference are first updated, and then each of the free images is non-rigidly registered onto the new reference. However, the existing knotpoints should require only minimal refinement.

Selection of New Reference Image At each iteration the algorithm considers selecting a different image as the reference. The choice of reference image is important because all other images are warped onto it, so ideally it should be as representative of the set as possible. For this reason, the candidate reference is selected to be the image that minimises the sum of pairwise distances to all of the other images. However, the reference is only updated if doing so does not degrade the set of objective function scores (i.e., does not damage the quality of the registration of the group as a whole).

The knotpoint positions on this candidate reference $Q_1^{(S')}$ are refined before deciding whether to adopt it as the new reference. This is done by perturbing the knotpoint positions (to move it out of possible local minima), and then re-optimising over the $Q_1^{(S')}$ using non-rigid registration to either another random image, or the current reference. We find that misalignment of an image can cause it to be erroneously selected as a candidate reference (see Fig. 1 for an example of this). The refinement step above gives us a chance to automatically correct this misalignment.

Data-driven Knotpoint Selection A set of knotpoints are defined on the reference image as the basis for the first pass through the algorithm, and this set of knotpoints is extended on each pass, based on the groupwise discrepancy between the free images and the current reference. The choice of knotpoint selection

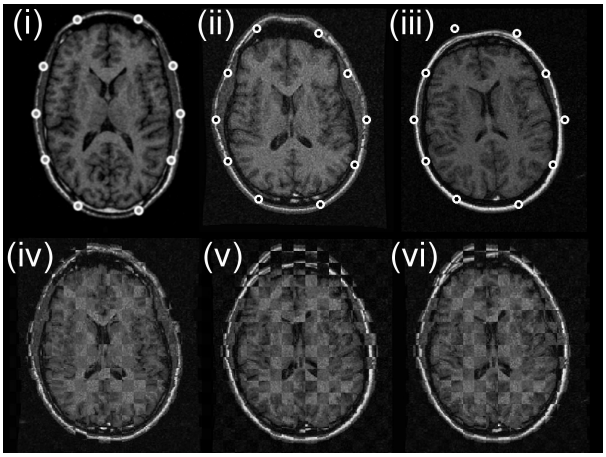

Fig. 1. Demonstration of refinement of knotpoints on candidate new reference. In plots (iv) to (vi), images are pulled back to the frame of the example free image (ii). **(i)** The initial reference image and automatically generated initial knotpoints. **(ii)** Example free image. **(iii)** Unrefined candidate reference. **(iv)** Interlaced example free images, showing relative alignment. **(v)** & **(vi)** Interlaced example free and candidate reference showing alignment before and after refinement.

method will depend on the particular image set being used. In the experiments described in section 4, the initial set of knotpoints were automatically generated to be spaced at equal angles around the centre of the image and placed on strong edges, which generally corresponded with the skull (Fig. 1). The new knotpoints that are added during each iteration of the algorithm are placed in regions where the quality of the registration, averaged across the group, is worst; that is, in the areas that display the greatest misalignment (Fig. 2). It should be noted that this method of groupwise non-rigid registration is *not* related to landmark-based methods of rigid or non-rigid registration (e.g., [13, 2]). Such landmark-based methods rely on finding significant *corresponding* landmarks on all images, whereas our method only uses landmark identification on a single image, as an aid to initialising the knotpoints on the first pass through the algorithm.

4 Experiments

To demonstrate the action of our algorithm, we took a set of 13 2D T1-weighted MR images of normal brains (see Fig. 4). The 227×227 pixel images were from 13 different individuals, and the slices chosen to show the anterior and posterior parts of the lateral ventricles. A reference image was chosen at random (see Fig. 1), and 10 initial knotpoints generated automatically, as described previously. The boundary circle was chosen as the circumcircle of this reference image. Each free image was affinely aligned to this reference, optimising over scaling, rotation and translation.

A first pass through the algorithm was then performed, using mutual information as the objective function; a subset of results are shown in Fig. 4. For this choice of objective function, the discrepancy is defined as follows. For each free image, the bins of the intensity histogram for that free image each define a set of target pixels. This set of pixels in the free image corresponds to a set of pixels in the reference, for which we construct a model of the intensity values across the set. The discrepancy for a particular reference pixel is then the probability of the occurrence of that pixel value, according to the intensity model of the reference set to which it belongs.

This small number of knotpoints on the skull aligns all 12 free images reasonably well, despite the large variations in skull shape across the set. There are slight misalignments for just 2 free images (2^{nd} and 5^{th} in Fig. 4) – these particular images do not have a large CSF-filled space at the front of the skull, and the registration has aligned the skulls from these images with the front of the brain of the reference. A new reference was chosen (image 5^{th}, see Fig. 1 (iii)). Note that the registration between free images (Fig. 1 (iv)) is considerably better than between the proposed reference and one of the same free images (Fig. 1 (v)) – this is because of the misalignment of image 5. The knotpoints for this candidate reference were then refined; however, even after refinement, the original reference still performed better in terms of the mean mutual information across the set (see Fig. 3), so was retained. Some additional knotpoints were chosen (see Fig. 2), and a second pass of the algorithm performed. The successive improvements of the mutual information across the group is shown in Fig. 3.

Fig. 2. The MI discrepancy image averaged across the set, with new knotpoints.

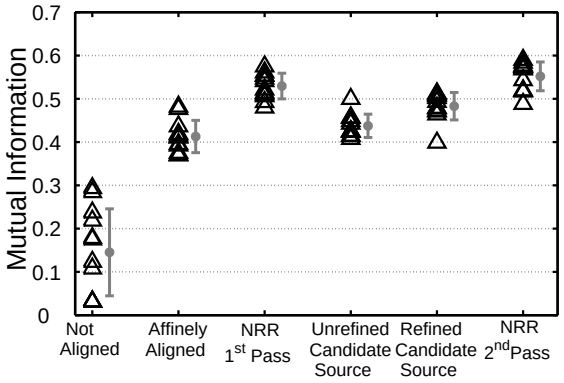

Fig. 3. MI across the set as the algorithm progresses. *Grey bars:* mean and standard deviation of the data.

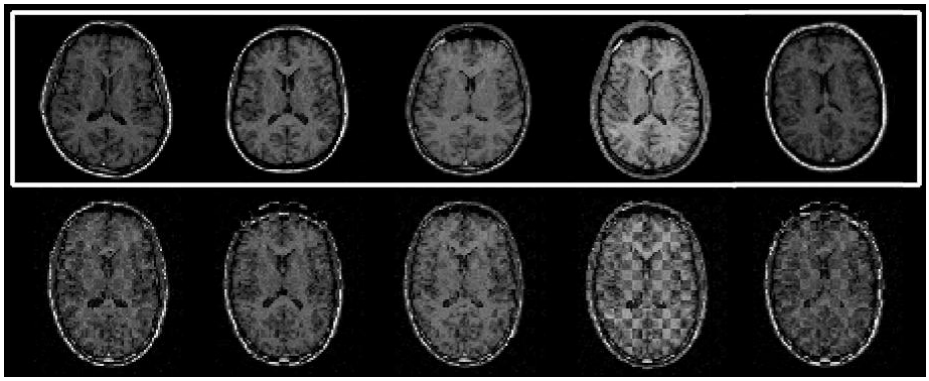

Fig. 4. First pass: non-rigid registration to the initial reference and knotpoints. Some of the (affinely aligned) free images are shown in the white box with, underneath, the interlaced images of the reference and the registered free images. The initial reference is in Fig. 1 (i).

5 Discussion and Conclusions

This paper has introduced an algorithm for non-rigid registration of groups of medical images. The algorithm is based on splines that interpolate between a set of knotpoints defined on a reference image; these knotpoints are transferred onto all of the free images, with their positions being optimised, so that there is a common low-dimensional representation of the group of images. Furthermore, the choice of reference image can be changed by the algorithm, based on the warp distances between all of the images in the set, and additional knotpoints are added at each iteration based on the groupwise discrepancy between the free images and the new reference in order to improve the registration.

The algorithm has been demonstrated on a set of 13 axial MR slices of the human brain. We have shown the ability of the algorithm to automatically correct misaligned examples, and to correctly register the group of images, despite the low-dimensionality (14 knotpoints) of the common representation. The experiments described were based on an implementation of the algorithm coded in MATLAB; a single pass through the algorithm took approximately 10 hours

on a 1.8 GHz P.C. The algorithm can be extended to 3D without difficulty (the description of the triharmonic CPS in 3D is given in [17]); however the computational costs will require a compiled implementation.

Acknowledgements This research was supported by the MIAS IRC project, EPSRC grant number GR/N14248/01.

References

1. T. Boggio. Sulle funzioni di green d'ordine m. *Rendiconti - Circolo Matematico di Palermo*, 20:97–135, 1905.
2. F. L. Bookstein. Principal Warps: Thin-Plate Splines and the Decomposition of Deformations. *IEEE PAMI*, 11(6):567–585, 1989.
3. M. Bro-Nielsen and C. Gramkow. Fast fluid registration of medical images. In *Proceedings of Visualization in Biomedical Computing (VBC)*, pages 267–276, 1996.
4. R. Chandrashekara, R. H. Mohiaddin, and D. Rueckert. Analysis of myocardial motion in tagged MR images using nonrigid image registration. In *Proceedings of 6th Medical Image Understanding and Analysis*, pages 1 – 4, 2002.
5. C. Chefd'Hotel, G. Hermosillo, and O. Faugeras. A variational approach to multimodal image matching. In *Proceedings of IEEE Workshop on Variational and Level Set Methods (VLSM'01)*, pages 21 – 28, 2001.
6. Rh. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor. 3D statistical shape models using direct optimisation of description length. In *Proceedings of ECCV*, Lecture Notes in Computer Science 2352, pages 3–20, 2002.
7. D. R. Forsey and R. H. Bartels. Hierarchical B-spline refinement. *ACM Transactions in Computer Graphics*, 22(4):205–212, 1988.
8. J. Gee, M. Reivich, and R. Bajcsy. Elastically deforming 3D atlas to match anatomical brain images. *Journal of Computer Assisted Tomography*, 17(2):225–236, 1993.
9. A. Guimond, J. Meunier, and J.-P. Thirion. Average brain models: A convergence study. Technical Report RR-3731, INRIA, Sophia Antipolis, 1999.
10. 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.
11. S. Marsland and C. Twining. Constructing data-driven optimal representations for iterative pairwise non-rigid registration. In *Biomedical Image Registration*, 2003.
12. C.R. Rohlfing, T. Maurer,Jr. Intensity-based non-rigid registration using adaptive multilevel free-form deformation with an incompressibility constraint. In *Proceedings of MICCAI*, Lecture Notes in Computer Science 2208, pages 111–119, 2001.
13. K. Rohr, H. S. Stiehl, R. Sprengel, T. M. Buzug, J. Weese, and M. H. Kuhn. Landmark-based elastic registration using approximating thin-plate splines. *IEEE Transactions on medical imaging*, 20(6):526–534, 2001.
14. 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*, Lecture notes in Computer Science 2208, pages 77–84, 2001.
15. D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes. Non-rigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.
16. C. Studholme, D. Hill, and D. Hawkes. Automated three-dimenensional registration of magnetic resonance and positron emission tomography by multiresolution optimisation of voxel similarity measures. *Medical Physics*, 24(1):25 – 35, 1997.
17. C. J. Twining and S. Marsland. Constructing diffeomorphic representations of non-rigid registrations of medical images. In *Proceedings of IPMI*, 2003.
18. C. J. Twining, S. Marsland, and C. J. Taylor. Measuring geodesic distances on the space of bounded diffeomorphims. In *Proceedings of BMVC'02*, 2002.