The Constrained Distance Transform: Interactive Atlas Registration With Large Deformations Through Constrained Distances

Bill Hill and Richard A. Baldock
MRC Human Genetics Unit
Edinburgh, EH4 2XU, UK
Bill.Hill@hgu.mrc.ac.uk
Richard.Baldock@hgu.mrc.ac.uk

Abstract

We describe a modification to distance based transforms used for non–linear registration, such as radial basis function transforms, which is both simple and novel. In the Constrained Distance Transform (CDT), distances are constrained to geodesics within the domain of the object being transformed. We show that the modified method is capable of producing large meaningful deformations for both an extreme synthetic case and for real biological image data. We further show that the method is suitable for interactive use on current workstations.

1 Introduction

Spatially mapped databases are becoming increasingly common in biomedical research. Some of these databases, such as EMAGE [1], contain volumetric models onto which spatially organised data are mapped through non–linear spatial transformations.

Two connected problems that are frequently encountered when defining these transformations are: Large deformations, such as those required for the pose of limbs or a curvature in the main body of an object, frequently fail to produce meaningful results; secondly, very different deformations are frequently required for two regions which are close in Euclidean space but distant in terms of the object. This last situation is illustrated in figure 1, in which despite the two points H and T being close in Euclidean space they are distant in terms of the object being transformed. It is clear that Euclidean geometry does not provide a good description of organisms with complex shapes and that a natural coordinate system might be more appropriate. A natural coordinate system can be defined as a "coordinate system which requires the least residual deformation to explain variability across individuals" [8]. Defining transformations which will uncurl the object or deform the head without deforming the tail are not possible using current methods, but it is just such transformations that are required to establish mappings between the assay and atlas model images shown later in figure 4.

A feature of the images shown in figure 4 is the variability of the assay images, which frequently do not correspond in either image values or their gradients. It is clear that, given the state of existing automatic registration algorithms, some interaction or expert guidance

is required to register the assay images to the atlas if the registration is to generate high quality mappings.

Existing methods which allow large deformations include articulated and fluid models. Martín-Fernández et al. [4] have developed articulated models for registering hand radiographs. In these models a skeleton composed of articulated rods is registered using landmarks at the ends of the rods and displacements away from the rods are interpolated using weighted combination of affine transforms. These models allow large deformations but do not appear to be a natural model for objects that do not have rigid components and in such cases a large number of articulated components might be required to achieve a reasonable transformation. Fluid models based on solving viscoelastic systems have been developed for registration with large deformations. While these models are often successful in producing large deformations, they are computationally expensive, non–interactive and do not always give biologically meaningful results: "Hands and brains do not in general deform like honey"[5].

Radial basis function (RBF) transforms are well suited to interactive use and are frequently used successfully for defining transforms with small deformations. But unfortunately these transforms do not behave well for large deformations or for large deformation gradients and even for smooth transformations they may not give diffeomorphic mappings [5].

In this paper we propose a simple modification to RBF transforms which allows them to be used interactively to produce high quality mappings for large deformations on common desktop workstations.

Figure 1: From left to right: Volume rendering showing points in the head (H) and tail (T), diagram showing the path for the constrained distance between H and T, distance transform showing the distance of all points in the domain from H. In the distance transform image the distance values have been normalised to the range 0 (black) to 255 (white) for clarity.

2 Radial Basis Function Transforms

Because of the difficulty in computing warp transformations automatically, RBF transforms based on manually or semi–manually defined landmark points are frequently used in biomedical applications. These manually defined landmarks can be combined with automatically generated landmarks through various weighting schemes. It is also possible

to combine landmark based registration with intensity based registration [3]. Landmark points give the set of displacements between the defined coordinates of the point pairs, but to achieve a warp transformation it is necessary to have a displacement function that is defined for all points within the domain of the source object. RBF transforms, in which the displacement at any point is given by the sum of weighted functions of the radial distances from the landmark points, are frequently used in this role.

Given a source object $O_s(x,y)$ and a target object $O_t(u,v)$, the displacements Δu and Δv defined by

$$\Delta u = u - x \tag{1}$$

$$\Delta v = v - y \tag{2}$$

can be computed using a RBF transform with the general form

$$\Delta u = P_u(x,y) + \sum_{i=1}^{i=N} \lambda_i b(r_i) \tag{3}$$

$$\Delta v = P_v(x,y) + \sum_{i=1}^{i=N} \mu_i b(r_i) \tag{4}$$

where P_u and P_v are two first order polynomials (the affine component), N is the number of landmark points, λ_i and μ_i are the basis function coefficients and b is the chosen basis function. In these equations r_i is the distance of some point $p(x,y)$ from landmark point l_i.

Two frequently used basis functions are the thin plate spline (TPS) and the multi-quadric (MQ), which in 2D have the form

$$b_{TPS}(r) = r^2 ln r^2 \tag{5}$$

$$b_{MQ}(r) = \sqrt{r^2 + \delta^2}. \tag{6}$$

The value of the MQ regularisation parameter δ is chosen so as to balance the smooth variation in displacements against exact landmark point displacements. This parameter is application dependent, but in practise it can be computed using a simple algorithm [2].

From a set of landmark points a system of linear equations may be written using the coefficients of the polynomial (P_u or P_v) and the RBF (λ_i or μ_i). These equations, known as the design equations, may then be solved using a linear system solver such as singular value decomposition. In practise it is often beneficial to rescale the parameters to reduce the condition number of the design matrix [2].

3 Distance Transforms

A distance transform maps a set of points within some domain of interest to their distance from a reference domain. Most distance transform algorithms assume a single rectangular (or at least convex) domain of interest, but in this work no such restriction is imposed and both the domains may be non–convex. In this more general case, the distances are evaluated along geodesics that are restricted to the domain of interest.

The distance transform algorithm used is based on a region growing algorithm described by Piper and Granum [7], but which has been implemented using efficient morphological primitives based on interval coding [6]. Given two domains, Ω_r the reference

domain and Ω_f the domain specifying the region of interest, a domain with a thin shell Ω_i is iteratively expanded from it's initial domain corresponding to the reference domain Ω_r. At each iteration Ω_i is dilated and clipped by it's intersection with Ω_f until Ω_i becomes the null domain $\emptyset$. At each iteration the current distance is recorded in a value table which covers the domain Ω_f. This is shown in algorithm 1. To approximate Euclidean distances, the domains are rescaled using a scale parameter and dilation is performed using a circular structuring element with a radius equal to the scale parameter. This rescaling is relatively efficient because the morphological primitives are based on interval processing and the number of intervals increases with a domain's linear dimension not it's area.

Algorithm 1 Distance Transform.

Require: Ω_f, Ω_r

$c \leftarrow$ initial connectivity
$s \leftarrow 0$
$V_d \leftarrow$ new integer value table with domain Ω_f
$\Omega_i \leftarrow \Omega_r$
while $\Omega_i \neq \emptyset$ **do**
  $s \leftarrow s + 1$
  $\Omega_p \leftarrow \Omega_i$
  $\Omega_i \leftarrow \mathit{diffdom}(\Omega_p, \mathit{dilation}(\Omega_p, c))$
  $\Omega_d \leftarrow \Omega_i \cap \Omega_f$
  $V_d|\Omega_d \leftarrow s$
end while

4 Constrained Distance Transforms

RBF transformations are relatively easy to implement and are well suited for interactive use. But in practise there are many situations in which they fail to produce useful transformations, such as where large deformation gradients are required.

In this paper we propose that a solution to the problem of interactively registering images with large deformations can be found through reconsidering the distances used by the RBFs (and other distance based transforms) and constraining the paths along which the distances are evaluated to be geodesics that lie within the domain of the source object. With regard to figure 1, it is clear that the head and tail are close in Euclidean space and that any landmarks which create a large displacement for the head will result in a similar displacement for the tail. However if the path along which this distance is measured is constrained to a geodesic within the object's domain, then the head is far from the tail despite it's being close in Euclidean space. By using distances constrained to the source object's domain, large displacements applied to the head may be effectively independent of those applied to the tail. All that is required is to substitute constrained distances in place of all Euclidean distances in the RBFs. We call this transformation the Constrained Distance Transform (CDT).

When evaluating the displacement of the source domain at some point p_i, using RBFs defined using landmarks $\{l_j\}$, then all distances between p_i and l_j must be computed for all i and j. Because constrained distances are considerably more expensive to compute

than Euclidean distances efficiency gains can be made by caching the the distance transforms, with one distance transform cached per landmark pair. Using this approach the cached distance transforms may be created and destroyed incrementally in step with the interactively defined landmark points.

In this work the transformation of both object domains and values are accomplished using non–regular triangulated conforming meshes. The use of conforming meshes simplifies image re–sampling in situations, such as limb articulation, which would otherwise give rise to large mesh element deformations requiring complex re–meshing schemes. Objects are transformed by first defining a mesh covering the source domain. This mesh then has a displacement, computed by evaluation of the radial basis function transform, associated with each of it's nodes. The mesh is first used to forward transform the source object's domain and then, if the source object has values associated with it, a new value table is created for the transformed domain and the new values are set using a sweep–line algorithm which scans through the transformed domain setting it's values from the source object via the current mesh element's inverse transform.

5 Results

An initial 2D implementation of the CDT has been evaluated using both synthetic and real biological images. In all cases a multiquadric RBF was used and the approximate Euclidean distance transform scale factor was set to 10.0.

A pair of synthetic images was generated corresponding to a large deformation problem which is known to be problematic for unconstrained RBFs [5]. These images are shown in figure 2. The 'C' image has a domain that is an incomplete annulus, with radii of 100 and 200 pixels and a gap sector of 30°. The 'I' image is an axis aligned rectangle with dimensions 100×800 pixels. Eighteen landmark pairs were computed for the C and I images during their generation and these were used to transform both the C to the I and the I to the C images using a CDT. The resulting images clearly show that CDTs are capable of registering objects requiring large deformations. that are problematic for unconstrained RBFs.

An assay image was selected from the EMAGE database in which the head and tail are close in Euclidean space. The image was segmented from its background using interactive grey value thresholding combined with connected component labelling. Landmark points were then evenly distributed around the segmented image, with all the landmarks having zero displacements except for those in the tail which had a displacement set away from the head. Figure 3 shows this image together with the landmark points and the displacements set. This figure also shows the results of applying a CDT and an unconstrained RBF transform to the image. When a CDT was applied, the head was not subjected to any significant deformation despite the large deformation experienced by the tail. However the unconstrained RBF transform resulted in large deformations to both the head and the tail. In both cases the same landmarks and displacements were used.

Figure 4 shows a set of eight typical 2D assay images from a set of submissions to the EMAGE database. This figure also shows a 2D projection of the corresponding 3D atlas model. These images display the typical variation that is seen in both the pose of the organisms and in their image values. Unconstrained RBF transforms are unable to map the assay images onto the atlas image satisfactorily. Attempts to do so result in folds

of the source image and a confused image in which multiple foreground regions may be blended with background. All assay images were first segmented from their background using interactive grey value thresholding combined with connected component labelling and any extraneous parts (such as the allantois) were removed. The segmented images were then mapped to the 2D atlas image interactively using CDTs, with the landmarks being placed at plausible correspondences. Figure 5 shows the successfully mapped assay images together with the target atlas image. For these images the number of landmark pairs used varied between 31 and 45.

In it's current implementation, the CDT computation time is dominated by the time taken to compute the distance transforms, with this accounting for approximately 99% of the total computation time. To compute the approximate distance transform, in which all distances from a single landmark are computed within the domain of the C image (outer diameter 400 pixels), takes 2.1s in the worst case on a 2.2GHz Opteron CPU using the approximate Euclidean distance transform with a scale factor of 10.0. This is the time delay between a user entering a landmark pair and seeing the displaced image computer.

6 Conclusions

RBF transforms defined using landmarks points are easily implemented and are frequently used for interactive non–rigid registration, however it is well known that these functions often fail to produce satisfactory warps for large deformations. A novel modification of the method has been presented which makes use of constrained distances. The CDT has been shown to produce meaningful results for a synthetic dataset that is known to be problematic for the unmodified method. It has also been shown to be capable of decoupling the displacements of regions that are close in Euclidean space, provided that an adequate segmentation can be achieved. Further, the CDT has been shown to produce satisfactory mappings, requiring large deformations, between biological assay images and a standard atlas model.

The computational cost of the approximate distance transforms currently restricts the size of image that can be transformed interactively. With one distance transform being computed for each landmark pair, the source is limited to 2D objects of a few hundred pixels in each dimension, such as the C image. Although it is possible to use a lower scaling factor in the approximate Euclidean distance transform, this results in a lower quality mapping which may be unacceptable. One method of making the method interactive, even for 3D data, is to map the atlas to the assay image, using landmarks fixed in the atlas image and pre–computed distance transforms. The resulting transform can then easily be inverted with mesh based implementations.

Future developments will include the addition of compact RBFs and the extension of the work to 3D transformations.

References

[1] R. Baldock, J. Bard, A. Burger, N. Burton, J. Christiansen, G. Feng, B. Hill, D. Houghton, M. Kaufman, J. RAo, J. Sharpe, A. Ross, P. Stevenson, S. Venkataraman, A. Waterhouse, Y. Yang, and D.R. Davidson. EMAP and EMAGE: A framework for understanding spatially organised data. *Neuroinformatics*, 1:309–325, 2003.

[2] R.A. Baldock and B. Hill. *Image Warping and Spatial Data mapping*, chapter 9. Oxford Uiversity Press, Oxford, UK, 2000. Editors: Baldock, R. and Graham, J.

[3] B. Fischer and J. Modersitzki. Combination of automatic non-rigid and landmark based registration: the best of both worlds. In M. Sonka and J. M. Fitzpatrick, editors, *Medical Imaging 2003: Image Processing*, pages 1037–1048. SPIE Press, 2003. Proceedings of the SPIE 5032.

[4] M. Á. Martín-Fernández, E. Muñoz Moreno, M. Martín-Fernández, and C. Alberola-López. Articulated registration: Elastic registration based on a wire-model. In J. M. Fitzpatrick and J. M. Reinhardt, editors, *Medical Imaging 2005: Image Processing*, pages 182–191, San Diego, CA, USA, February 12–17 2005. SPIE Press. Proceedings of the SPIE 5747.

[5] Jan Modersitzki. *Numerical Methods for Image registration*. Oxford University Press, Oxford, UK, 2004.

[6] J. Piper and D. Rutovitz. Data structures for image processing in a C language and Unix environment. *Pattern Recognition Letters*, 3:119–129, 1985.

[7] Jim Piper and Erik Granum. Computing distance transformations in convex and non-convex domains. *Pattern Recognition*, 20(6):599–615, 1987.

[8] Daniel Rueckert, Alejandro F. Frangi, and Julia A. Schnabel. Automatic construction of 3D statistical deformation models of the brain using non-rigid registration. *IEEE Transactions on Medical Imaging*, 22(8):1014–1025, 2003.

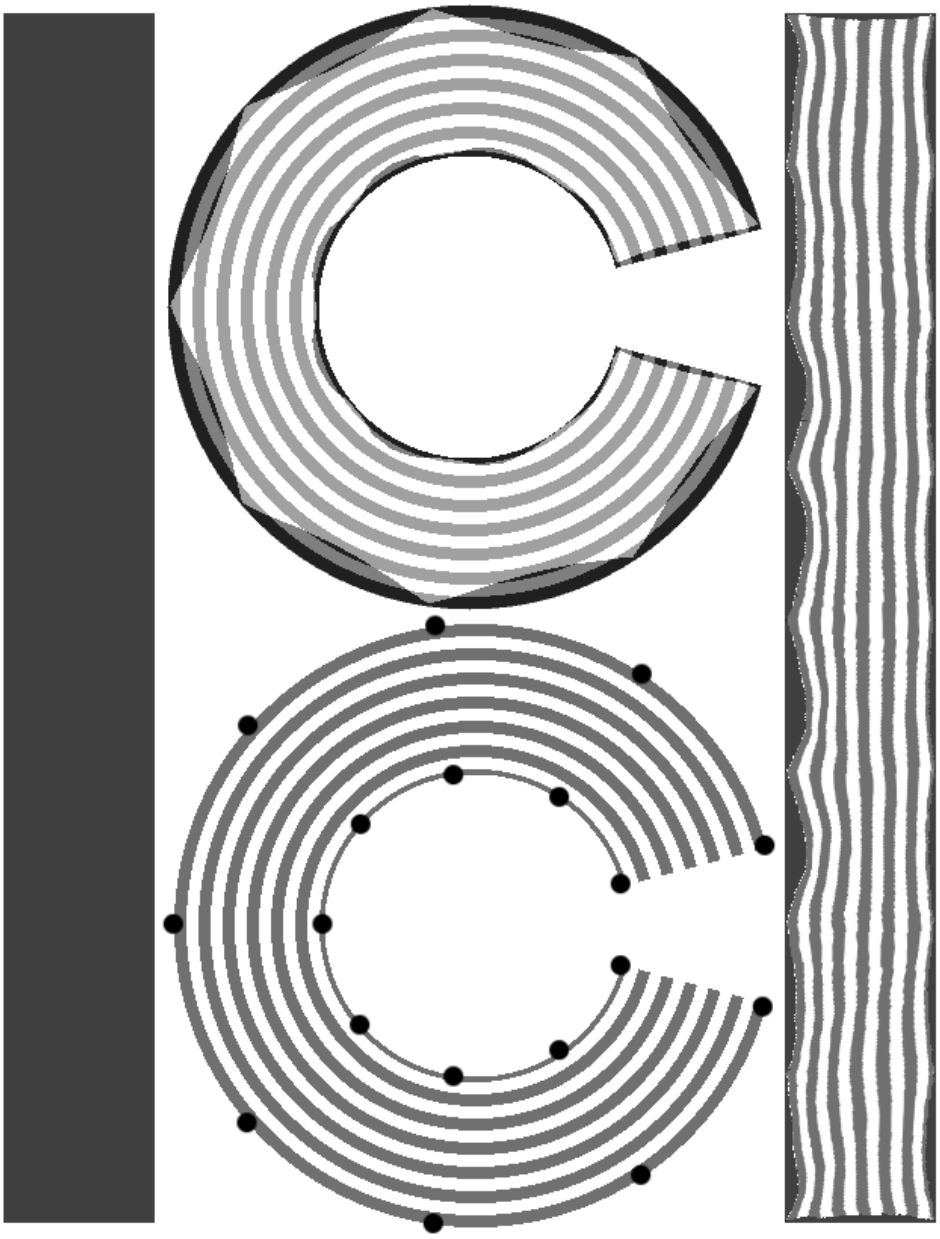

Figure 2: C and I dataset. Counter clockwise from the left: I object, C object showing the positions of the landmark points, overlay of the C object warped to the I, overlay of the I object warped to the C.

Figure 3: From left to right: Source image (from EMAGE) showing fixed points and displacements, source image transformed using a CDT and source image transformed using an unconstrained RBF transformation.

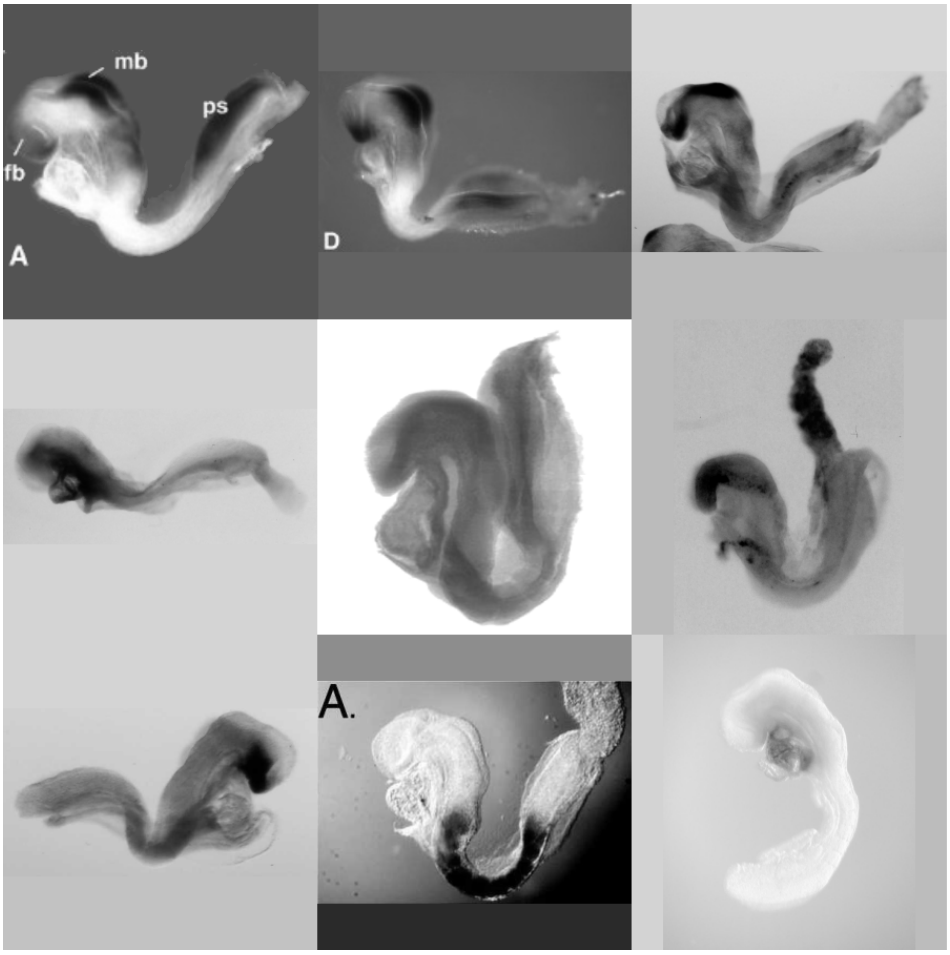

Figure 4: EMAGE TS12 dataset: Outer images are assay images to be mapped to an atlas image (centre).

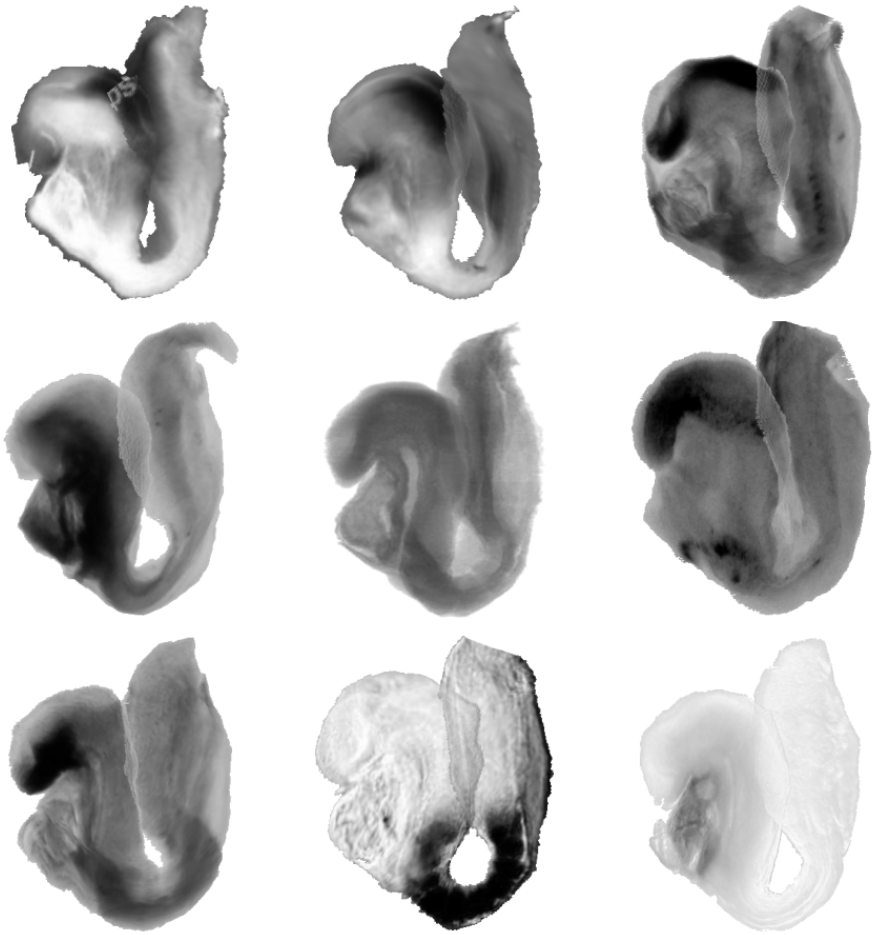

Figure 5: EMAGE TS12 dataset: Outer images are the registered assay images mapped to an atlas image (centre).