Non-rigid Registration by Geometry-Constrained Diffusion

Per Rønsholt Andresen^{1,2} and Mads Nielsen^{2,3}

¹ Department of Mathematical Modelling, Technical University of Denmark

pra@imm.dtu.dk

<http://www.imm.dtu.dk/~pra>

² 3D-Lab, School of Dentistry, University of Copenhagen, Denmark

<http://www.lab3d.odont.ku.dk>

³ IT-University of Copenhagen, Denmark

<http://www.it.edu>

Abstract. Assume that only partial knowledge about a non-rigid registration is given so that certain points, curves, or surfaces in one 3D image map to certain points, curves, or surfaces in another 3D image. We are facing the aperture problem because along the curves and surfaces, point correspondences are not given. We will advocate the viewpoint that the aperture and the 3D interpolation problem may be solved *simultaneously* by finding the *simplest* displacement field. This is obtained by a geometry-constrained diffusion which yields the simplest displacement field in a precise sense. The point registration obtained may be used for growth modelling, shape statistics, or kinematic interpolation. The algorithm applies to geometrical objects of any dimensionality. We may thus keep any number of fiducial points, curves, and/or surfaces fixed while finding the simplest registration. Examples of inferred point correspondences in a longitudinal growth study of the mandible are given.

Keywords: Aperture-problem, automatic landmark detection, diffusion, kD-tree, non-rigid registration, simplest displacement field, homology.

1 Introduction

In a registration, we wish to establish the spatial correspondence of points in two images. Correspondence is defined through the concept of homology except in pathological cases. In general, homologous points will, dependent on the medical task, reflect similar anatomy, functionality, or geometry, etc. In this paper, we assume that homologous *objects* have been defined *a priori*. Therefore, we seek an automatic method for establishing *point* correspondences based on *object* correspondences. Pursuing this, we presume that: 1) The optimal registration is a mapping between homologous points, 2) The underlying biological process is smooth and homologous points do not “change place” i.e., the ordering of the anatomical structures is preserved. Formally: the registration field must not fold or be torn apart. It is then a homeomorphism.

Fig. 1. The images show schematically how the diffusion algorithm works on the deformation field. The Cartesian components of the initial deformation field (arrows in the left image) are Gaussian smoothed. Some of the links have now diverged from the surface (middle image) and must be projected back on to the surface (right image). The fold (the two crossing arrows) is removed by repeating the steps until the field does not change.

In other words, within the objects, a solution to the generalized *aperture problem* must be chosen. In this paper, we introduce the concept of *geometry-constrained diffusion* for solving the interpolation and aperture problems simultaneously.

When performing shape statistics or analyzing (longitudinal) shape development, the tools from shape statistics (see e.g., [4]) require point matches. That is, to perform a statistical analysis of the variation of shapes we must identify homologous *points* on the shape samples. When having only a few landmarks the registration may be performed manually, but for thousands of points it becomes tedious and practically impossible. In many cases punctual landmarks are hard to establish in images, and the process requires considerable prior anatomical knowledge.

Automated methods using geometrical features such as crest-lines [9] are powerful, but do not provide a *dense* field, and may give problems in regions where shape features change topology so that correct matching is not possible. We propose using *geometry-constrained diffusion* for inferring the locally simplest non-rigid object registration.

The result of *geometry-constrained diffusion* is a dense, continuous, invertible displacement field (a homeomorphism). Many fields may fulfill the geometrical constraints given by the objects. The diffusion process gradually simplifies an initial registration field. In general, diffusion is a gradient ascent in entropy. That is, locally it changes the registration field so as to remove its structure as fast as possible. An unconstrained diffusion in this way leads to an affine registration. The geometry-constrained diffusion also simplifies the registration field as fast as possible, but is limited locally so as to preserve the object mappings (see Fig. 1).

In section 3, the theory of geometry-constrained diffusion is summarized. Section 4 describes the implementation. Examples of the simplification of an initial crest-line-based non-rigid registration are shown in section 5.

2 Related Work

In the literature, many algorithms for non-rigid registration exist. In this paper, we address the equally important problem of measuring the complexity of the geometrical deformation in a non-rigid registration. This measure may be introduced either by having only a finite number of semi-local low parameter registrations, or a viscous fluid or elasticity constraint, or a deformation energy of which the thin-plate spline energy is the canonical example (see [6] for a survey). Feldmar and Ayache's approach[5] resembles ours the most.

Feldmar and Ayache[5] perform a surface registration based on a distance measure including local geometrical properties of the surfaces. The surface registration is a collection of local affine registrations. The parameters of these registrations are spatially blurred so as to construct a smoothly varying registration. A difference to our approach is that we do not make a collection of local affine frames, but a global registration field. Secondly, and most importantly, we do not exploit any metric properties of the surfaces, but look for a globally simple registration field. This also creates a tendency to match points of similar geometry since the field otherwise must be more complex.

In principle, the geometry-constrained diffusion could also have been formulated as a geometry-constrained gradient descent in displacement energy [3]. Hence, we here present a general technique for handling under-determined geometrical constraints in conjunction with variational approaches for non-rigid registration.

3 Geometry-Constrained Diffusion

A registration field may be diffused simply by diffusing the Cartesian components independently. The geometry-constrained diffusion is constructed such that it preserves certain fiducial mappings during the diffusion. Assume that the identification of some fiducial points, curves and/or surfaces is given *a priori*. In order to handle this partial geometrical knowledge in the general non-rigid registration problem, we propose geometry-constrained diffusion which in a precise sense simplifies the displacement field while preserving fiducial points, curves, and/or surfaces.

Given two images $I_1 : \mathbb{R}^3 \mapsto \mathbb{R}$ and $I_2 : \mathbb{R}^3 \mapsto \mathbb{R}$, we define the registration field $R : \mathbb{R}^3 \mapsto \mathbb{R}^3$. Along the same line we define the displacement field $D : \mathbb{R}^3 \mapsto \mathbb{R}^3$ such that $R(x) = x + D(x)$. We may then define:

Definition 1 (Displacement diffusion). *The diffusion of a displacement field $D : \mathbb{R}^3 \mapsto \mathbb{R}^3$ is an independent diffusion in each of its Cartesian components:*

$$\partial_t D = \Delta D$$

where the Laplacian, $\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$, is applied independently in the x -, y -, and z -component of D .

The only difference between the registration and displacement field is the addition of a linear term. This term does not influence the diffusion so that registration diffusion is identical to the displacement diffusion.

This vector-valued diffusion has some obvious and important symmetries:

Proposition 1. *The displacement diffusion is invariant with respect to similarity transforms of any of the source or target images.*

Proof. The translational part of the similarity transform only adds a constant to the displacement field, and the diffusion is invariant to this. The displacement $y = D(x) + x$ is (up to a translation) similarly transformed such that $x' = s_1 R_1 x$ and $y' = s_2 R_2 y$ where R_1 and R_2 are 3×3 rotation matrices. Under $s_1 R_1$ the displacement is mapped to $D_1(x') = D(s_1^{-1} R_1^{-1} x') - x' + s_1^{-1} R_1^{-1} x'$. Applying $s_2 R_2$ also we find

$$D'(x') = s_2 R_2 [D(s_1^{-1} R_1^{-1} x') - x' + s_1^{-1} R_1^{-1} x']$$

The latter terms leave the diffusion unaltered since they only add terms of first order, and the diffusion depends only on terms of second order. Since the diffusion is linear, it is invariant to $s_2 R_2$. By re-mapping t the diffusion is known to be independent of similarity transforms of the base manifold. •

Applying the displacement diffusion without further constraints, it reaches a steady state which is an affine registration. This is easily seen since only linear functions are in the null-space of the diffusion equation.

In the case where the same geometrical structures have been identified in both images we wish to make certain that the diffusion of the displacement field reflects these structures. Assume that a surface $S_1 : \mathbb{R}^2 \mapsto \mathbb{R}^3$ in the source image is known to map on to the surface $S_2 : \mathbb{R}^2 \mapsto \mathbb{R}^3$ in the target image. We thus define

Definition 2 (Surface-constrained diffusion). *The surface constrained diffusion of $D : \mathbb{R}^3 \mapsto \mathbb{R}^3$ mapping $S_1 : \mathbb{R}^2 \mapsto \mathbb{R}^3$ onto $S_2 : \mathbb{R}^2 \mapsto \mathbb{R}^3$ is given by*

$$\partial_t D = \begin{cases} \Delta D - n_{S_2} \frac{n_{S_2} \cdot \Delta D}{\|n_{S_2}\|^2} & \text{if } x \in S_1 \\ \Delta D & \text{if } x \notin S_1 \end{cases}$$

where n_{S_2} is the unit surface normal of $S_2(D(x) + x)$.

This corresponds to solving the heat flow equation with certain boundary conditions. In this case, however, we do not keep the solution fixed at the surface, but allow points to travel along the surface. This is a dual approach to the geometry-driven curve and surface diffusion by Olver, Sapiro, and Tannenbaum [8] and others. We keep only the tangential part of the diffusion along the surface whereas they diffuse the geometry of the surface maintaining only the normal flow. The surface normal n_{S_2} may simply be obtained as a length normalization of $n_{S_1} + J n_{S_1}$ where J is the Jacobean of D . In this way the formulation is no longer explicitly dependent on S_2 . That is, given an initial (guess of the) displacement field and a surface in this source image to be preserved under

diffusion, we may still apply the above equation without explicitly referencing to S_2 .

Curve constraints and point constraints can be handled in a similar manner. For the curve problem, we project onto a curve by only taking the part of the diffusion which is along the curve tangent. Point constraints simply disregard the diffusion at these points. The three types of geometry-constrained diffusions may be combined in any fashion as long as the boundary conditions (the matches) do not contradict one another.

We make the following proposition:

Proposition 2 (Similarity Invariance (II)). *The geometry-constrained diffusion is invariant to similarity transforms of the source or target image.*

Proof. We have already shown that the unconstrained diffusion is similarity invariant. Both the surface normal and the curve tangent are also invariant under the similarity transform. •

We will conjecture that the geometry-constrained diffusion removes any fold in the initial displacement. This means that, the steady state solution to the geometry-constrained displacement diffusion creates an invertible mapping.

Conjecture 1 (Invertibility). A geometry-constrained diffused displacement field induces a one-to-one mapping of $\mathbb{R}^3$.

The steady state displacement field will be a homeomorphism assuming the above invertibility-conjecture is valid since the constrained diffused displacements are continuous. It will also be smooth except on the constrained objects where it will generally not be differentiable across object boundaries, but will be differentiable along smooth objects.

It is evident that the scheme is not symmetric in the images. This is due to the change in local metric by the non-linear displacement field. This makes the ordering of the two images important. It is, however, not obvious (to us) that the steady states will differ.

The geometry-constrained diffusion can be implemented simply by applying an numerical scheme for solving a space and time discretized version of the diffusion. It is well known that the diffusion equation is solved by Gaussian convolution. That is, an unconstrained diffusion can be updated an arbitrarily long time-step, by applying a Gaussian of appropriate size. The geometry-constrained diffusion cannot be solved directly in this manner due to the constrains. In general, the finite time step diffusion (Gaussian convolution) makes the displaced source surface diverge from the target surface, so that it must be back-projected to the target surface. The back-projection may be performed to the closest point on the target surface (see Fig. 1). In this way, the algorithm resembles the iterative closest point algorithm [2,10] for rigid registrations.

Fig. 2. Flow diagram for the diffusion algorithm. See section 4 for details.

4 Implementation

A time and space discretized solution the the geometry-constrained diffusion may be obtained by iterative Gaussian convolution and back-projecting the constrained surfaces.

The crux of the algorithm then becomes (see Fig. 2 for a flow chart):

- 1. **Initial displacement.** Construct an initial guess of the displacement field.
- 2. **Diffusion step.** Convolve the displacement field with a Gaussian kernel.
- 3. **Deform source.** Deform the source surface with respect to the displacement field.
- 4. **Matching (Projection onto the target surface).** For all points on the deformed surface: Find the closest point on the target surface.
- 5. **Update displacement field.** For all points on the deformed surface: Change the displacements according to the match.
- 6. **Convergence.** Is the displacement field stable? If not, go to 2.

Some of the steps are explained in greater detail below.

Fig. 3. Iso-surface and crest-lines for a 3 (left) and 56 (right) month old mandible. The mandibles are Gaussian smoothed ($\sigma = 3\text{mm}$) in order to capture the higher scale features. The dimensions of the left and right mandibles are ($H \times W \times L$) $18 \times 57 \times 53\text{mm}$ and $31 \times 79 \times 79\text{mm}$, respectively. Surfaces are translucent.

4.1 Diffusion Step

We use the normalized Gaussian convolution [7]. For each of the Cartesian components of the displacement field, a Gaussian weighted average is constructed and divided with the sum of the weights. The standard deviation of the Gaussian σ is the only parameter in the numerical scheme (see section 4.4).

Fig. 4. Result of running the diffusion algorithm ($\sigma = 2\text{mm}$) on the displacement field. Deformation of the 56 month old mandible to the 3 month old mandible (see Fig. 3). The surface and wire-frame of the deformed surface are shown to the left and right respectively. The initial displacement, one iteration with the diffusion algorithm, and the last iteration are shown from top to bottom, respectively. The foldings are a result of the imperfect initial registration (extremal-mesh registration extended to the whole surface by Gaussian regularization as in [1]). The final result is almost perfect, but some folds still exist, owing to the discretization of the surface and displacement field.

4.2 Matching

As in [10] we use a 3D-tree for finding the closest point on the target surface. As reference points on the triangulated target surface we use the center of mass

(CM point) for each triangle. The three corners of the triangles are used for calculating a plane. Using also the surface normals, we construct the following algorithm for finding the closest point: First, find the closest CM point using the Kd-tree. Secondly, calculate the closest point on the surface as the intersection of the corresponding triangle-plane and the line given by the deformed point and the normal at the CM point.

4.3 Convergence

The diffusion is stopped when

$$\sum_{p_i} \|D_n(p_i) - D_{n-1}(p_i)\|^2 < \epsilon, \quad (1)$$

where p_i is the points on the source surface, D_n is the displacement in the n th-iteration, and ϵ is a user-chosen parameter. Alternatively, a fixed number of iterations could be chosen. 5-10 iterations is normally enough.

4.4 Choice of Time Step σ

The Gaussian kernel size, σ , is the only parameter in the diffusion algorithm. It determines the time discretization step. A too large value of σ , may tear apart the surface since we diffuse too far before back-projecting. This problem occurs in regions of high surface curvature. A too small value of σ also gives problems since we have chosen a fast but imprecise back-projection algorithm. The error in the back-projection introduces some artificial “bumps” in the path along which we diffuse. This may be overcome by a more precise back-projection algorithm or in practice by choosing σ sufficiently large (see Fig. 4-6). Identical solution are obtained for an interval of σ 's.

In practice, we choose a small σ and increase it on the fly if folds persist.

5 Results

The method has been applied for registration of 31 mandibles from 6 different patient in a longitudinal growth study of the mandible. One mandible is chosen as the reference mandible. In order to propagate the landmarks, all mandibles are registered with the reference mandible and geometry-constrained diffusion is applied. The reference mandible is shown in Fig. 3-right. Fig. 3-left displays the target surface for all the subsequent figures, except for Fig. 6, which shows an example where the diffusion algorithm gives an erroneous result. The prior estimate of the displacement field is obtained by crest-line matching [1]. See Fig 7 for match between two sets of crest-lines.

As seen in Fig. 4 (top images) the initial deformation contains folds. Applying the diffusion algorithm removes almost all the folds, but some persist. By increasing σ (see section 4.4), these are removed (Fig. 5). As seen in Fig. 6, too large a value of σ will eventually tear apart the surface.

Fig. 5. Converged diffusion algorithm with a high value of σ ($\sigma = 10\text{mm}$). The surface and wire-frame of the deformed surface are shown to the left and right respectively. We have forced the displacement field to be more smooth, by increasing σ .

Fig. 6. The deformation vectors are moved too far away from the surface (The value of σ is too high) resulting in a wrong projection back onto the surface.

Very convincingly, Fig. 8 shows that the crest-lines are useful anatomical landmarks but only in areas where their topology stays fixed. Teeth eruption changes the crest-line topology of the mandible. We see two lines before teeth eruption on top of the mandible (Fig. 3 - left image) but only one after teeth eruption (Fig. 3 - right image). A pure (crest-) line matching algorithm is not able to handle such changes. Introducing the diffusion algorithm, the single crest-line (the green line on top of the mandible in Fig. 8) is able to perform correctly - i.e., be registered in between the two other lines (the two red lines on top of the mandible in the same figure) as seen in Fig. 8-right.

The same phenomenon is seen on the bottom of the mandible. A single line on the young mandible is split in two on the older mandible.

6 Conclusion

In the present paper we have proposed an algorithm for finding the simplest displacement field, which is conjectured to be a homeomorphism (1-1 continuous mapping).

Fig. 7. Match (lines in black) between the two sets of crest-lines (before applying the diffusion algorithm). The crest-lines in red and green are from the mandibles shown in Fig. 3. Only every 11th link is shown for visual clarity. We see that the matches to a very good extent connect homologous points.

The geometry-constrained diffusion in this paper serves to simplify the non-rigid registration of surface models. The result is a much smoother displacement field. Volume registration is achieved by having more than one surface. It turns out that the algorithm itself is also very simple.

In theory, the method is parameter free, but implementations include parameters of space- and time-discretization and convergence threshold.

We are currently using the method for registering a longitudinal growth study of the mandible in order to extract more than 14000 homologous points which again are used for inference of the growth. In that study, applying the geometry-constrained diffusion results in a very significant increase in the explained variance by the growth model.

Acknowledgments

The work was supported by the Danish Technical Research Council, grant number 9600452 to Per Rønsholt Andresen. The authors also thank Sven Kreiborg (School of Dentistry, University of Copenhagen, Denmark) and Jeffrey L. Marsh (Plastic and Reconstructive Department for Pediatric Plastic Surgery, Washington University School of Medicine at St. Louis Children's Hospital, St. Louis, Missouri, USA) for the CT data. Also thanks to Bjarne K. Ersbøll (Technical University of Denmark) and Andy Dobrzeniecki (3D-Lab, Denmark) for comments on the manuscript. The Visualization Toolkit (<http://www.kitware.com>) was used for the visualizations.

Fig. 8. Left and right images show the deformed (in green) and the original (in red) crest-lines before and after applying the diffusion algorithm ($\sigma = 2\text{mm}$), respectively. In the initial registration crest-lines are registered with crest-lines. Where the topology does not change and away from umbilic points we see (almost) no movement of the green crest-lines. Teeth eruption changes the topology on “top of the surface” (see Fig. 3) therefore the green crest-lines move.

References

1. P. R. Andresen, M. Nielsen, and S. Kreiborg. 4D shape-preserving modelling of bone growth. In W. M. Wells, A. Colchester, and S. Delp, editors, *Medical Image Computing and Computer-Assisted Intervention (MICCAI'98)*, volume 1496 of *Lecture Notes in Computer Science*, pages 710–719, Cambridge, MA, USA, 1998. Springer. Electronic version: <http://www.imm.dtu.dk/~pra>. 539, 540
2. P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239 – 255, 1992. 537
3. F. L. Bookstein. Landmark methods for forms without landmarks: localizing group differences in outline shape. In *Proceedings of the IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 279–89, 1996. 535
4. F. L. Bookstein. Shape and the information in medical images: A decade of the morphometric synthesis. *Computer Vision and Image Understanding*, 66(2):97–118, 1997. 534
5. J. Feldmar and N. Ayache. Rigid, affine and locally affine registration of free-form surfaces. *International Journal of Computer Vision*, 18(2):99–119, 1996. 535
6. H. Lester and S. R. Arridge. A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32:129–149, 1999. 535
7. M. Nielsen and P. R. Andresen. Feature displacement interpolation. In *IEEE 1998 International Conference on Image Processing (ICIP'98)*, pages 208–12, 1998. Electronic version: <http://www.imm.dtu.dk/~pra>. 539
8. P. J. Olver, G. Sapiro, and A. Tannenbaum. Invariant geometric evolutions of surfaces and volumetric smoothing. *SIAM Journal on Applied Mathematics*, 57(1):176–194, 1997. 536
9. J.-P. Thirion. The extremal mesh and the understanding of 3D surfaces. *International Journal of Computer Vision*, 19(2):115–128, 1996. 534
10. Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *International Journal of Computer Vision*, 13(2):147–176, 1994. 537, 539