Med cal Image Analys s (2000) volume ?, number ?, pp 101–109

Non-rigid Registration by Geometry-Constrained Diffusion

Per Rønsholt Andresen[1*] and Mads Nielsen[2]

[1]Department of Mathematical Modelling, Technical University of Denmark, Denmark
[2]IT University of Copenhagen, Denmark

Abstract
Assume that only partial knowledge about a non-rigid registration is given: certain points, curves, or surfaces in one 3D image are known to map to certain points, curves, or surfaces in another 3D image. In trying to identify the non-rigid registration field, we face a generalized aperture problem since 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 in a precise sense yields the simplest displacement field. The point registration obtained may be used for segmentation, growth modeling, shape analysis, 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 synthetic example and a longitudinal growth study of the human mandible are given.

Keywords: Aperture-problem, automatic landmark detection, simplest displacement field, homology.

Received ?; revised ?; accepted ?

1. Introduction

In a registration, we wish to establish the spatial correspondence of points in two (3D) images. When performing shape statistics or analyzing (longitudinal) shape development, the tools from shape statistics (e.g., (Bookstein, 1991; Bookstein, 1997b; Dryden and Mardia, 1998)) 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.

A homologous point or a (semi-) landmark is a point the correspond across all the cases of a data set under a reasonable model of deformation from their common mean (Bookstein, 1991).

In this paper, we assume that homologous *objects* have been defined *a priori*. By a homologous *object* we understand the same anatomical object (e.g. the mandible). 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 are preserved. Formally: the registration field must not fold or be torn apart. It is then a homeomorphism.

We introduce the concept of *geometry-constrained diffusion* for solving the registration problem given an initial deformation field. Examples show that the method is capable of giving the correct point-to-point correspondence when the initial displacement field only comprises the relation between the objects, i.e. given an intial (guess of the) displacement field, the algorithm fully automatically establish the correct point correspondence.

The result of geometry-constrained diffusion is a dense, continuous, invertible displacement field (a homeomorphism). Many fields may fulfill the geometrical constraints

*Corresponding author
(e-mail: pra@imm.dtu.dk)

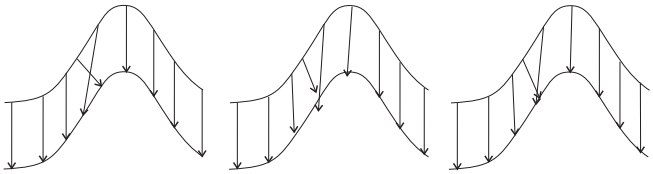

Figure 1. The images show schematically how the geometry-constrained diffusion algorithm works on the deformation field. The Cartesian components of the initial deformation field (arrows in the left image) are smoothed by Gaussian kernels. Some of the links have now diverged from the surface (middle 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.

given by the objects. The diffusion process gradually simplifies an initial registration field. In general, diffusion is a gradient ascent in entropy, i.e. the optimal way to move matter so as to even out the distribution of matter. That is, locally a diffusion will change the registration field so as to remove its structure as fast as possible. An unconstrained diffusion of a registration field leads, as we will show later, 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 (Figure 1).

Often, it is very easy to establish the initial object correspondence: for the cube an initial random generated correspondence still gives the optimal result when using the algorithm, similar results are also obtained for the mandibular test-case (Section 5). This also means that the method not only regularizes an a priori displacement field, it is a registration method that is capable of giving the optimal displacement field based on the (very weak) assumption: that points on the source object must map to points on the target object.

In this paper, the objects are (closed) surfaces in 3D space, but the method can easily be generalized to objects of any dimension where one wish to improve (simplify) the displacement field.

The organization of the remainder of this text is as follows. Section 2 reviews related work. 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 (e.g. (Maintz and Viergever, 1998; Lester and Arridge, 1999) for surveys). The diffusion process introduced in this paper may be viewed as a gradient descend of a corresponding energy functional. In the geometry-constrained version, the relation is most obvious in terms of biased diffusion shown formally to correspond to a gradient descend of a regularizing functional (Nordström, 1990). Such a regularizing functional has been introduced for registration 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 ((Lester and Arridge, 1999) for a survey). Feldmar and Ayache's approach (Feldmar and Ayache, 1996) resembles ours the most.

Feldmar and Ayache (Feldmar and Ayache, 1996) 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, we make a global registration field. Second, 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 (Bookstein, 1997a). Hence, we here present a general technique for handling underdetermined geometrical constraints in conjunction with variational approaches for non-rigid registration.

Automated methods using geometrical features such as crest lines (Thirion, 1996; Thirion and Gourdon, 1995; Thirion and Gourdon, 1996; Subsol et al., 1998; Andresen et al., 1998) 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. Other automated methods using geometrical features have been evaluated in the literature (Hartkens et al., 1999), but this is outside the scope of the present paper.

3. Geometry-constrained Diffusion

A registration field may be diffused simply by diffusing the Cartesian components independently. An overview of linear and non-linear scalar and vector valued diffusion schemes may be found in (ter Haar Romeny, 1994). The geometry-

constrained diffusion is constructed such that it preserves certain fiducial mappings during the diffusion.

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:*

$$\partial_t D = \triangle D$$

where the Laplacian, $\triangle = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$, is applied independently on 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) similarity 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

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

The later terms (***MADS SKAL RETTE DETTE, DA DET IKKE ER KORREKT***) 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 exist 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 respects 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*

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

where n_{S_2} is the 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 an approach dual to the geometry-driven curve and surface diffusion by (Olver et al., 1997) 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. This corresponds in spirit to Bookstein's concept of *deficient coordinate* (Bookstein, 1991). The surface normal n_{S_2} may simply be obtained as $n_{S_1} + J n_{S_1}$ where J is the Jacobian 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 S_2. However, this is not the approach we have implemented (Section 4).

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

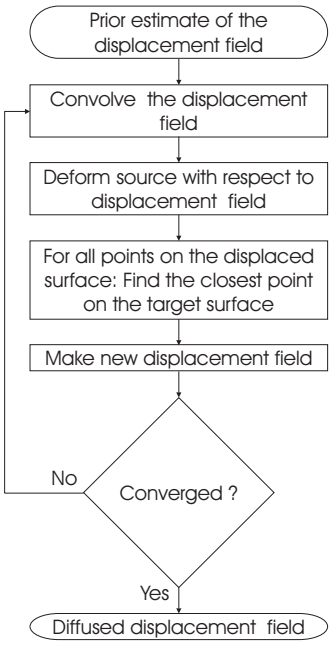

Figure 2. Flow diagram for the diffusion algorithm. Section 4 describes the details

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) whatever the steady states will differ or not.

The geometry-constrained diffusion can be implemented by applying an numerical scheme for solving a space and time discretized version of the diffusion. It is well known that the diffusion equation can be 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 constraints. 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 (Figure 1). In this way, the algorithm resembles the Iterative Closest Point (ICP) algorithm (Besl and McKay, 1992; Zhang, 1994) for rigid registrations.

4. Implementation

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

The crux of the algorithm then becomes (Figure 2 shows the 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.

4.1. Diffusion step

We use the normalized Gaussian convolution (Nielsen and Andresen, 1998). For each of the Cartesian components of the displacement field, a Gaussian weighted average is constructed and divided by the sum of the weights. The standard deviation of the Gaussian σ is the only parameter in the numerical scheme (Section 4.4).

4.2. Matching

As in (Zhang, 1994) 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.

We construct the following algorithm for finding the closest point: First, find the closest CM point using the Kd-tree (Preparata and Shamos, 1988) (,here $K = 3$). A list of length k containing the triangle itself (called the CM triangle) and the triangles neighbors are made. Neighbors are defined as triangles sharing at least one corner point with the CM triangle. The closest point for each triangle in the list is calculated as the shortest distance, $d_i, i \in [1, \cdots, k]$, between the deformed point and the triangle. The closest point is then found as the point having the shortest distance among d_i.

The proposed algorithm has the advantage that the deformed points can move continuously on the target surface.

4.3. Convergence

The diffusion is stopped when

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

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

4.4. Choice of Time Step σ

The Gaussian kernel size, σ, is the only parameter in the diffusion algorithm. σ determines the time discretization step. As in any numerical scheme for solving partial differential equations (where the analytical solution is unknown), we should not take too long time steps. A too large value of σ, may tear the surface apart since we diffuse too far before back-projecting. This problem arises first in regions of high surface curvature combined with a deformation field having too low resolution. Decreasing σ and/or increasing the resolution of the deformation field solves the problem. The only drawback of a small σ is that more iterations are needed before convergence is reached (two convolutions with a Gaussian of size σ equals one convolution with a Gaussian of size $\sqrt{2}\sigma$). See Figures 9-11.

5. Results

This section shows examples on synthetic and real images.

For both types of examples, a non-rigid registration algorithm based on crest-lines and Adaptive Gaussian Filtering gives an initial "guess" of the registration (Figure 4 and 7). *Adaptive Gaussian Filtering* (Nielsen and Andresen, 1998) calculates a dense deformation field from a sparse one and is defined as separate normalized Gaussian convolutions for each Cartesian coordinate where the kernel size in each point equals the square root of the Euclidean distance to the nearest feature (here the crest lines).

The registration of crest lines is done as in (Subsol et al., 1998) and (Andresen et al., 1998): 1) first calculating a rigid registration by moments (affine transformation) (Bajcsy and Kovacic, 1989); 2) then two first order polynomial deformations; 3) and two second order polynomial deformations (Brown, 1992); 4) lastly a totally non-rigid deformation (all (deformed) points are moved freely to the nearest point in the target set). The point to point correspondence in each of the four steps is made by the ICP algorithm. See (Glasbey and Mardia, 1998) for a more recent review of (2D) image warping methods.

From the sparse deformation field given by the crest lines, a dense deformation field is calculated by Adaptive Gaussian Filtering.

Having the initial dense deformation field, we use the geometry-constrained diffusion to gradually simplify the field. In this process the diffusion is constrained by the surface not the crest lines.

Notice that the initial registration could have been done by any method: one does not have to use crest lines as geometrical features for the initial registration (e.g. (Hartkens et al., 1999)). Therefore, in order to test the capabilities of the algorithm, random initial deformations fields was generated and used as the initial "guesses".

5.1. Cube Registration

This section shows how the algorithm behaves on a synthetic data set. In a 64^3 volume a 32^3 cube is generated with $1mm$ cubic resolution. The volume is Gauss smoothed ($\sigma = 4mm$) in order to compute the crest lines. A small and a large cube are extracted from the Gauss smoothed volume by the marching cubes algorithm (Lorensen and Cline, 1987) with iso-value 240 and 90, respectively (Figure 3). The two sets of crest lines and the registration between them are shown in Figure 4.

Figure 5 left image shows the initial deformation of the large cube using only the registration given by the crest lines and Adaptive Gaussian Filtering. The right image in Figure 5 shows how the deformed large cube looks when using the converged displacement field. It is seen that we have achieved the optimal registration, namely a scaling of the large cube to the size of the small cube (the wireframe of the large cube (Figure 3) is identical to the wireframe of the deformed large cube using the converged deformation field (Figure 5-right)).

In order to test the stability of the geometry-constrained diffusion we conducted three different experiments.

1. **Deformation of endpoints.** n random pairs of endpoints of the converged deformation field (field from Figure 5-right) were swapped. The large cube consist of 6774 points (vertices). We swapped up to 60000 pairs (this means that some points are swapped more than one time) in order to make a graph showing the errors between the final field and the optimal field (the converged field). However, all the fields converge to the field obtained with no points swapped, making all the errors zero. A video shows how the iterations proceed for $n = 60000$. The video is found at `http://www.imm.dtu.dk/~pra/MedIA`.

2. **Random initialization.** All the endpoints of the deformation field were uniformly distributed randomly all over the volume ($[0;63]^3$ - both cubes have center of mass in (32,32,32)) and then projected to the closest

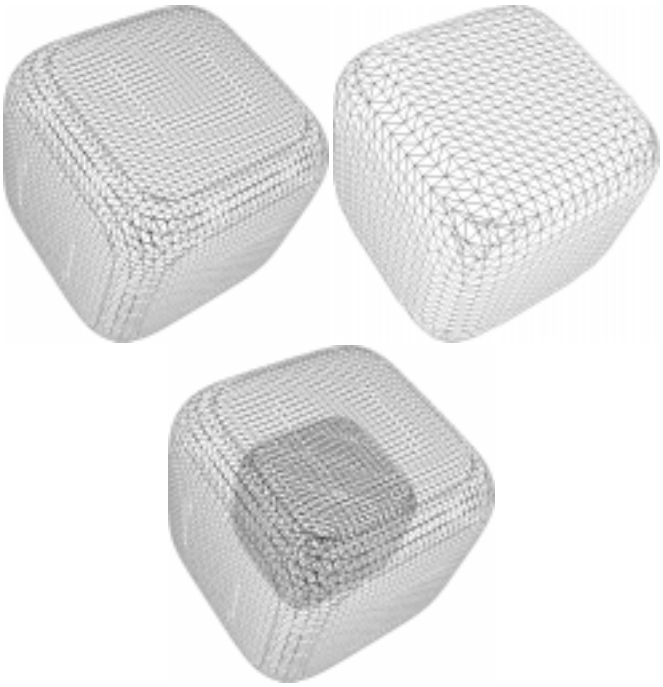

Figure 3. The wireframe of the large (top left) and small (top right) cube. Below, both cubes visualized in the same screen in order to compare their relative size. See Section 5.1 for the generation of the cubes.

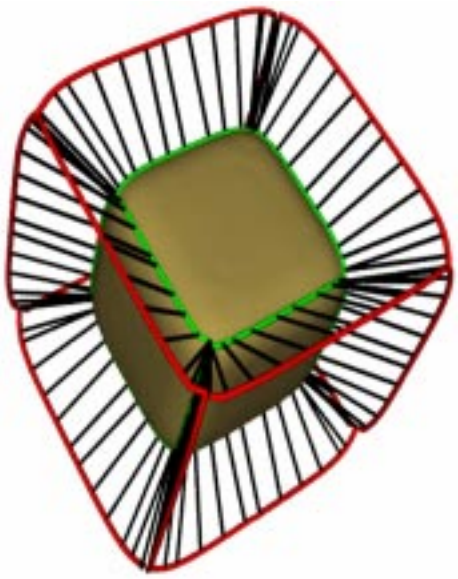

Figure 4. The red and green lines are the crest lines on the large and small cube, respectively. Links between the two set of crest lines are shown as black lines. For visual clarity only every third link is shown.

points on the cube. The field converges again. A video is found at the same web page.

3. **Singular initialization.** All the endpoints of the deformation field were given the value (0,0,0) and then projected to the closest point on the cube. The field does not converge to the initial field. This is also expected as the deformed cube would have to "tear itself apart" in order to encircle the target cube. A video is found at the same web page.

These three experiments show that the method is very stable. The value $\sigma = 2$ was used for the diffusion in all three experiments. In the second example ("Random initialization") the method converges after approximately 290, 20, and 3 iterations for $\sigma = 2$, 5, and 10, respectively (Videos are found on the web page). We are using Fast Fourier Transform (FFT) when doing the Gaussian filtering (instead of a convolution in the spatial domain), therefore this uses as much as 47 CPU seconds on a SGI Octane single 300MHz R12000 MIPS CPU. The matching ("Projection onto the target surface") uses 13 CPU seconds.

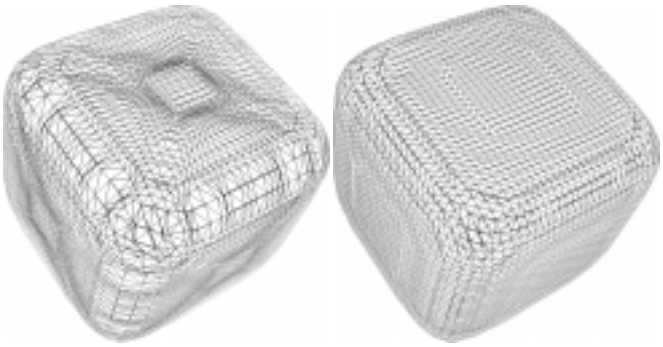

Figure 5. The sparse deformation field defined by the links (Figure 4) is extended to the whole surface by Adaptive Gaussian Filtering. Using this initial dense deformation field the deformation of the large cube to the small is seen as the left image. The right image shows the deformed surface when the diffusion has converged. Notice that the right wireframe (the converged deformation) is identical to the wireframe of the large cube (Figure 3-top left): only a scaling has been achieved.

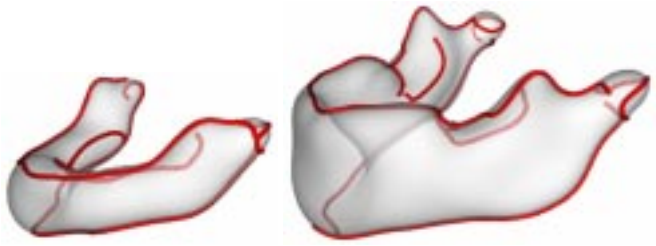

Figure 6. Iso-surface and crest lines for a 3 (left) and 56 (right) month old mandible. The mandibles are Gaussian smoothed ($\sigma = 3mm$) in order to capture the higher scale features. The dimensions of the left and right mandibles are (H$\times$W$\times$L) 18$\times$57$\times$53mm and 31$\times$79$\times$79mm, respectively. Scannings have $0.5mm$ cubic resolution. Surfaces are translucent.

Figure 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 Figure 6. We see that the matches very satisfactorily connect homologous points. Only every 11th link is shown for visual clarity.

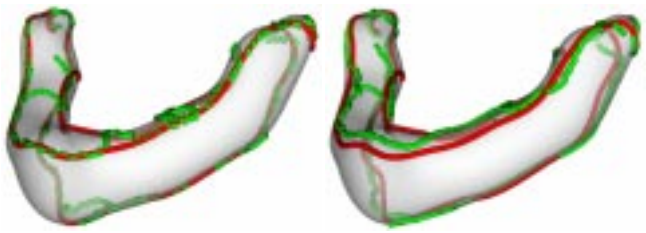

Figure 8. Left and right images show the deformed (in green) and the original (in red) crest lines before and after applying the diffusion algorithm ($\sigma = 2mm$), respectively. In the initial registration crest lines are registered with crest lines. Where the topology does not change and away from umbilic points (for which the curvature is the same in all directions e.g. at the condyles) we see (almost) no movement of the green crest lines. Erupting teeth change the topology on "top of the surface" (Figure 6) therefore the green crest lines move.

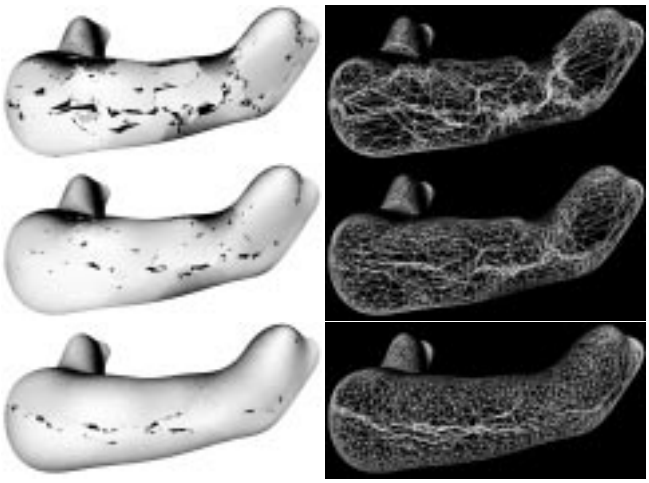

Figure 9. Result of running the diffusion algorithm ($\sigma = 2mm$) on the displacement field. Deformation of the 56 month old mandible to the 3 month old mandible (Figure 6). The surface and wire-frame of the deformed surface are shown to the left and right, respectively. The initial displacement, 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 (Figure 7) extended to the whole surface by Gaussian regularization as in (Andresen et al., 1998) - see also text in Section 5). The final result is almost perfect, but some folds still exist, owing to the discretization of the displacement field. A video is found at `http://www.imm.dtu.dk/~pra/MedIA.`

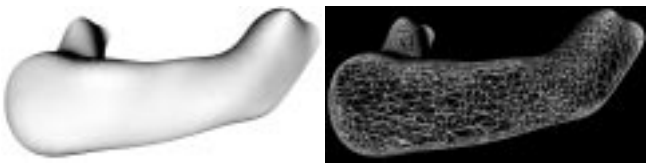

Figure 10. Converged diffusion algorithm with a high value of σ ($\sigma = 10mm$). The surface and wire-frame of the deformed surface are shown to the left and right, respectively. We have forced the displacement field to be more smooth, by increasing σ. See also caption in Figure 9.

Figure 11. 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. The resulting field is torn apart in the high curvature regions near the condyles and the teeth. See also text in Section 5.2.

5.2. Mandible Registration

The method has been applied for registration of 31 mandibles from 6 different patient in a longitudinal growth study of the mandible (Andresen et al., 2000; Andresen, 1999). One mandible is chosen as the reference mandible for the whole data set. In order to propagate the landmarks, all $31-1$ mandibles are registered with the reference mandible using the same registration technique as described in Section 5 and geometry-constrained diffusion is applied. The reference mandible is shown in Figure 6-right. Figure 6-left displays the target surface for all the figures involving mandibles (Figures 7-10), except for Figure 11, which shows an example where the diffusion algorithm, on purpose, is forced to give an erroneous result. The prior estimate of the displacement field is obtained by crest line matching as described in Section 5. Figure 7 shows the match between two sets of crest lines.

As seen in Figure 9 (top images) the initial deformation contains folds. Applying the diffusion algorithm removes almost all the folds, but some persist. By increasing σ (Section 4.4), these are removed (Figure 10). As seen in Figure 11, too large a value of σ will eventually tear apart the surface.

Very convincingly, Figure 8 shows that the crest lines are useful anatomical landmarks but only in areas where their topology stays fixed. Erupting teeth change the crest line topology of the mandible. We see two lines before tooth eruption on top of the mandible (Figure 6 - left image) but only one after tooth eruption (Figure 6 - 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 Figure 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 Figure 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.

One experiment was conducted for the mandible:

- **Deformation of endpoints**. The experiment "Deformation of endpoints" (See Section 5.1 for details) was also conducted for the mandible. 60000 pairs of endpoints were swapped (the reference mandible consists of 9087 vertices) and $\sigma = 8$. The video found at the web page (`http://www.imm.dtu.dk/~pra/MedIA`) shows how the field converges to the same optimal field as in Figure 10.

It is noteworthy that geometric-constrained diffusion is capable of finding the optimal field for such a complex shape as the mandible with random distributed matches[a].

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). The method has been presented for surfaces in 3D, but is easily generalized to any dimension.

The geometry-constrained diffusion in this paper serves to simplify the non-rigid registration of surface models but have also been demonstrated to work as a "pure" registration method. 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 (Andresen et al., 2000; Andresen, 1999). 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

[a]we allow ourself to call 60000 swapped endpoint (compared with 9087 vertices) for random distributed matches.

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), Andy Dobrzeniecki (3D-Lab, Denmark) and the reviewers for comments on the manuscript. The Visualization Toolkit (`http://www.kitware.com`) was used for the visualizations.

References

Andresen, P. R. (1999). *Surface-bounded Growth Modeling Applied to Human Mandibles*. PhD thesis, Technical University of Denmark, Lyngby, Denmark.

Andresen, P. R., Bookstein, F. L., Conradsen, K., Ersbøll, B. K., Marsh, J., and Kreiborg, S. (2000). Surface-bounded growth modeling applied to human mandibles. *IEEE Transactions on Medical Imaging*. Accepted for publication.

Andresen, P. R., Nielsen, M., and Kreiborg, S. (1998). 4D shape-preserving modelling of bone growth. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI'98*, volume 1496 of *Lecture Notes in Computer Science*, pages 710–719. Electronic version: http://www.imm.dtu.dk/∼pra.

Bajcsy, R. and Kovacic, S. (1989). Multiresolution elastic matching. *Computer vision, Graphics and Image Procession*, 46(1):1–21.

Besl, P. J. and McKay, N. D. (1992). A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239 – 255.

Bookstein, F. L. (1991). *Morphometric Tools for Landmark Data: Geometry and Biology*. Cambridge University Press.

Bookstein, F. L. (1997a). Landmarks methods for forms without landmarks: morphometrics of group differences in outline shape. *Medical Image Analysis*, 1(3):225–243.

Bookstein, F. L. (1997b). Shape and the information in medical images: a decade of the morphometric synthesis. *Computer Vision and Image Understanding*, 66(2):97–118.

Brown, L. G. (1992). A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376.

Dryden, I. L. and Mardia, K. V. (1998). *Statistical Shape Analysis*. John Wiley & Son.

Feldmar, J. and Ayache, N. (1996). Rigid, affine and locally affine registration of free-form surfaces. *International Journal of Computer Vision*, 18(2):99–119.

Glasbey, C. A. and Mardia, K. V. (1998). A review of image warping methods. *Journal of Applied Statistics*, 25:155–171.

Hartkens, T., Rohr, K., and Stiehl, H. (1999). Performance of 3D differential operators for the detection of anatomical point landmarks in MR and CT images. In Hanson, M., editor, *Proc. SPIE's International Symposium on Medical Imaging, Image Processing*, pages 32–43. SPIE-Int. Soc. Opt. Eng.

Lester, H. and Arridge, S. R. (1999). A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32:129–149.

Lorensen, W. E. and Cline, H. E. (1987). Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics*, 21(3):163–169.

Maintz, J. B. A. and Viergever, M. A. (1998). A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36.

Nielsen, M. and Andresen, P. R. (1998). Feature displacement interpolation. In *IEEE 1998 International Conference on Image Processing (ICIP'98)*, pages 208–212. Electronic version: http://www.imm.dtu.dk/∼pra.

Nordström, N. (1990). Biased anisotropic diffusion – a unified regularization and diffusion approach to edge detection. *Image and Vision Computing*, 8(11):318–332.

Olver, P. J., Sapiro, G., and Tannenbaum, A. (1997). Invariant geometric evolutions of surfaces and volumetric smoothing. *SIAM Journal on Applied Mathematics*, 57(1):176–194.

Preparata, F. P. and Shamos, M. I. (1988). *Computational Geometry. An Introduction. Corrected and Expanded Second Printing*. Springer-Verlag.

Subsol, G., Thirion, J.-P., and Ayache, N. (1998). A general scheme for automatically building 3D morphometric anatomical atlases: Application to a skull atlas. *Medical Image Analysis*, 2(1):37–60.

ter Haar Romeny, B. M., editor (1994). *Geometry-Driven Diffusion in Computer Vision*, volume 1 of *Computational Imaging and Vision*. Kluwer Academic Publishers.

Thirion, J.-P. (1996). The extremal mesh and the understanding of 3D surfaces. *International Journal of Computer Vision*, 19(2):115–128.

Thirion, J.-P. and Gourdon, A. (1995). Computing the differential characteristics of isointensity surfaces. *Computer Vision and Image Understanding*, 61:190–202. Electronic version: http://www.inria.fr/RRRT/RR-1881.html.

Thirion, J.-P. and Gourdon, A. (1996). The 3D marching lines algorithm. *Graphical Models Image Processing*, 58:503–509. Electronic version: http://www.inria.fr/RRRT/RR-1881.html.

Zhang, Z. (1994). Iterative point matching for registration of free-form curves and surfaces. *International Journal of Computer Vision*, 13(2):147–176.