Eurographics Workshop on 3D Object Retrieval (2011)
H. Laga, T. Schreck, A. Ferreira, A. Godil, I. Pratikakis, and R. Veltkamp (Editors)

Refining shape correspondence for similar objects using strain

L. Phan[1], A.K. Knutsen[1], P.V. Bayly[1], S. Rugonyi[2] and C. Grimm[1]

[1]Washington University in St. Louis, USA [2]Oregon Health and Science University, USA

Abstract

Several applications - for example, study of biological tissue movement and organ growth - require shape correspondence with a physical basis, especially for shapes or regions lacking distinctive features. For this purpose, we propose the adaptation of mechanical strain, a well-established physical measure for deformation, to the problem of constructing shape correspondence and measuring similarity between non-rigid shapes. In this paper, we demonstrate how to calculate strain for a 2D surface embedded in 3D. We then adjust the correspondence between two surfaces so that the strain varies smoothly across the deformed surface (by minimizing the change in strain). The final strain on the deformed surface can be used as a measure of shape similarity.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling —Physically based modeling

1. Introduction

With the advances in imaging techniques, such as microscopy, MRI, CT, and PET scanning, the amount of 3D data available has grown significantly. Considerable research effort has been concentrated on the problem of viewing or reconstructing 3D shapes from raw 3D data. However, recently the research focus has shifted to understanding, characterizing, and retrieving such 3D shapes from a database. To achieve these goals, an important task is to match shapes, construct a correspondence between shapes and quantify the change in shape based on that correspondence.

Both the parameterization and correspondence literature look at the problem of mapping one surface to another, with or without features. Given a set of corresponding features, either user-provided or automatically identified, there have been many approaches to establishing correspondence between surfaces. Some common approaches are maps that preserves angle, area, or distance [KS04][SPR06]. However, propagating the correspondence over flat regions between features, or regions without distinctive features, is open to interpretation.

To provide a physical basis to shape correspondence, we propose the use of strain as a deformation metric. Strain relates the forces on a material to measurable changes in its physical shape. Strain is well defined for n-D to n-D deformations. In this paper, we show how to define strain between two 2D surfaces embedded in 3D, for both analytical and discrete (e.g., mesh) surfaces of similar objects.

For correspondence extension, we use strain in a similar manner to the harmonic relaxation typically applied in this problem, except that we focus on minimizing the change in strain, not the strain itself. As a simple example, consider a rubber band stretched between two pins. One would expect the strain to be evenly distributed along the rubber band, not just accumulated in certain areas. Our proposed relaxation achieves that. We also show how mechanical strain can be applied to the problems of quantifying shape similarity and analysing shape deformation.

Establishing shape correspondence is the key to quantify and analyze shape change. This has many important real-life applications such as using imagery to track the development of biological tissue and internal organs. We test our approach by applying it to image data that tracks the motion of cardiac tissue in the outflow tract (OFT) of an embryonic chick heart.

Contributions

1. Computation of strain when source and destination shapes do not have the same mesh topology
2. Phrasing the problem of minimizing the variation in

DOI: 10.2312/3DOR/3DOR11/017-024

strain as a sparse matrix least-squares solve (Implicit integration)
3. Computing correspondence for cardiac tissue in motion for the study of cardiac morphology and function

2. Previous work

There are several approaches to establishing correspondence between surfaces, with parameterization [HPS08] being a special case of matching a surface to a domain. Previous energy-minimization approaches usually define energy functions consisting of components that, while allowing the vertices to slide on the first surface, pull them to the closest point on the second surface [HPS08]. In order to guide the matching process, the user must identify a sufficient number of corresponding feature points on both surfaces. Such energy functions only work for shapes that are close enough, otherwise the solver will converge to a poor local minimum [ACP03].

In cases where corresponding features are provided, or can be detected automatically, most of the existing methods have some form of relaxation for areas which are not directly constrained by feature lines. The elegant work of Miller et al. [BMTY05] defines a flow field in parameter space which takes feature lines to feature lines, and defines the flow in the remaining part of the parameter space so as to prevent fold-overs. Shi et al. [STD*07] use PDEs defined on level-sets to evolve one surface to the next. They use harmonic energy minimization to propagate the direction field at the feature lines across the rest of the surface. Van Essen et al. [EDD*01] use ad-hoc planar warps to align an initial conformal parameterization with desired feature lines. Gu et al. [WGC*08] also use conformal mapping, but provide landmark matching by explicitly cutting the mesh and using the features as boundaries; this prevents fold-overs. [GHDS03] on the other hand use stretch and bending energies to calculate the deformation of a mesh given the force acting upon it, which is a totally different problem.

Curvature is one of the geometric properties that can be used to establish surface correspondence for areas without distinctive features. However, curvature provides no input in areas with no strong features. Even in the presence of ridges, curvature provides matching across the ridge but not along it. Moreover, curvature is non-linearly affected by scale - this is particularly problematic in cases where the shape is undergoing an overall size change. Finally, curvature calculations tend to be noisy, which may affect any gradient-descent type algorithm.

Previous approaches in the literature discussed above use some deformation metric which works well for their task, but is not directly related to the underlying physics of the deforming material. To illustrate this, consider two common deformation metrics: the conformal metric and the area metric. Minimizing deformation using the conformal metric enforces uniform stretch in all directions. The shape of the local neighbourhood is preserved regardless of the change in size (a circle on one surface should be mapped to a circle on the other). Applying the area metric, on the other hand, disregards the direction of stretch and preserves the area allowing changes in shape (a circle can be mapped to an ellipse with arbitrary orientation but having the same area). Neither of these metrics (used independently) can capture the physical deformation, making it necessary to use heuristics (such as a suitably weighted sum) for combining them.

In this work, we use mechanical strain as a measure that truly captures the physical deformation. Although closely related to stretch, strain differs in that it captures both the direction and magnitude of deformation. We start with the mechanical strain calculated from an initial correspondence to measure the deformation of one shape to another. Then, by minimizing the change in strain across the deformed surface, we arrive at a final correspondence that yields the most meaningful deformation given the physical change from one shape to another. In particular, we reduce strain that arises primarily due to *parameterization*, and not fundamental changes in the shape.

Filas et al. [FKBT08] have shown how to calculate strain on a 2D surface using surface fitting. They break the problem into one polynomial for the height field and one (2D to 2D) for the parameterization, then extract the vectors and calculate a 3D matrix, from which they derive the strain values. This method is similar to (and is the inspiration for) our analytical calculation; we differ in that we use a far simpler polynomial fitting approach and derive the 2D tensor as well as the principal strain values.

In their recent work [KCG*10], Knutsen et al. construct the correspondence between two surfaces by minimizing an energy function that consists of two energy terms: the total strain due to distortion and an "error energy" due to mismatch between corresponding features. The strain relaxation is computed using the COMSOL software, which is based on a finite element method. The approach works for spherical surfaces only, does not prevent folding, is very slow, and requires many parameters to be specified. Instead of using a gradient descent approach, we relax the strain using an implicit forward solver. Our method is not only faster, but also allows extension to arbitrary topology, since it is implemented using a hierarchical grid-based approach [GSS99].

In this paper, we use strain to refine an initial correspondence between two similar objects. Such a correspondence can be constructed using standard approaches [SPR06]. The strain induced by actual deformation of the surface is spread out evenly across the surface as explained using the rubber band analogy (Section 1). However, the initial correspondence often induces an additional strain which leads to an uneven distribution of strain. Our goal is to modify the correspondence to relax the strain across the surface, effectively eliminating this unwanted variation in strain.

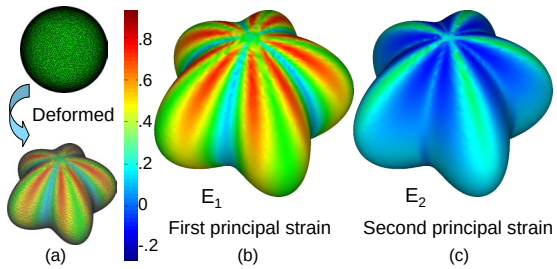

Figure 1: *(a) Strain induced when a sphere is deformed into a star-shape. Positive strain implies expansion, negative strain implies compression. (b) shows the first principal strain and (c) second principal strain.*

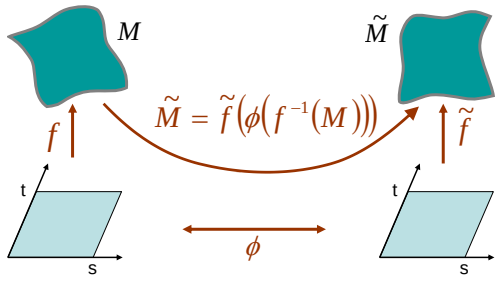

Figure 2: *Function f defines the undeformed surface M, while function $\tilde{f}$ defines the deformed surface $\tilde{M}$. Given a mapping ϕ between the two planar parameterizations, we can construct a bijective mapping between M and $\tilde{M}$.*

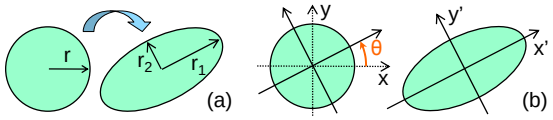

Figure 3: *(a) Magnitude of deformation is the stretch ratio between the radii r_i of the ellipse and the radius r of the circle, as shown in Eq. 2. (b) Direction of deformation is the rotation angle θ that aligns the circle's axes (x,y) with the ellipse's axes (x',y').*

We will present in Section 3 a background of mechanical strain. In Section 4, we show how to calculate strain for a 2D surface embedded in 3D using the initial correspondence. In Section 5 we will explain how we relax the strain using a grid-based implicit solver (see Fig. 5 for an overview of the algorithm flow). Finally, in Section 6 we will discuss the results of our experiments with this method.

3. Background on Strain

Strain is typically applied to a material, and as such, is defined for an n-dimensional (n-D) material embedded in n - D [Tab04]. In our setting, we only consider the strain induced on the infinitely thin 2D surface embedded in 3D, not the interior volume of the object. The surface can deform in three directions, although what we really want to capture is the change in the *geodesics* of the surface – this best captures the notion of measuring the change in the relative displacement of two points on the surface.

Figure 1 shows an example of strain on a sphere that has been deformed by alternately compressing and expanding it in a radial pattern. In the direction of compression, such as across the crease, the strain is negative. In the direction of expansion, such as across the ridge, the strain is positive. The two principal strains describe stretching in two orthogonal directions; for this example, the maximum strain largely captures the radial compression, the minimum the top-bottom stretch.

In the following paragraphs, we will give a geometric explanation of strain to provide insight into how we develop our calculation.

We define the 2D strain tensor, $F_{2D} \in \Re^2 \times \Re^2$, and from that the two principal strain measures ($E_1, E_2 \in \Re$) †, by looking at the transformation of a circle to an ellipse. We start with an infinitesimal circle around a point p in a 2D material. Now apply forces to that material causing the circle around p to deform to an ellipse (see Figure 3a). F_{2D} is therefore given by:

$$F_{2D} = R_\theta^T S R_\theta \tag{1}$$

where R_θ is a rotation matrix (where θ is the angle of rotation) that aligns the circle's x, y axes with the ellipse's major and minor axes (x',y'), and S is a scaling matrix that scales the circle to match the ellipse. Let r be the radius of the circle, and r_1 and r_2 be the length of the major and minor axes of the ellipse, respectively. Then, for $i = 1,2$, E_i is related to the change in the length of the axes [Tab04]:

$$E_i = \frac{1}{2}\left[\left(\frac{r_i}{r}\right)^2 - 1\right] \tag{2}$$

This implies that E_i will be zero if there is no strain. We can also extract E_i directly from F_{2D} using Singular Value Decomposition (SVD):

$$F_{2D} = U\Sigma V^* = R_{-\theta}\begin{bmatrix} E_1 & 0 \\ 0 & E_2 \end{bmatrix} R_\theta \tag{3}$$

Given F_{2D}, we define the strain energy as:

$$E = \sum_{i,j=1}^{2} F_{i,j}^2 \tag{4}$$

The two principal strains E_1, E_2 and the strain energy E

† E_1 and E_2 are the strains at the same point but in different directions. The first principal strain, E_1, is always larger than the second principal strain E_2.

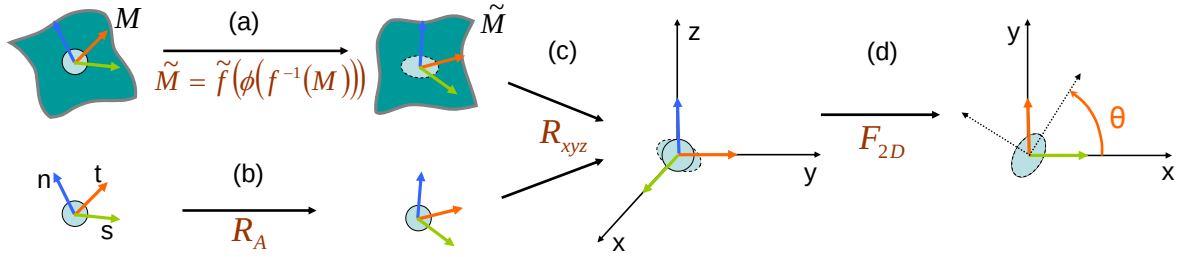

Figure 4: *(a) Bijective mapping between the reference surface M and the deformed surface $\tilde{M}$ (Fig. 2). (b) Rotation matrix R_a aligns the tangent frame at point p with that at point $\tilde{p}$. (c) The matrix R_{xyz} aligns the tangent frame at $\tilde{p}$ with the axes (x,y,z) such that the circle and the ellipse lie in the x-y plane as best as possible. (d) After dropping the z-dimension, the strain tensor F_{2D} deforms the circle into the ellipse.*

can be computed from F_{2D} as shown in equations 3 and 4. These measures of strain can be used to visualize the strain on the deformed surface. The next section explains in detail the calculation of the strain tensor F_{2D} on a mesh.

4. Strain calculation

In this section, we will describe how we calculate the strain tensor F_{2D}. Our discrete approach directly implements the geometric interpretation of strain (Section 3). Let M and $\tilde{M}$ be two surfaces defined by an embedding of the same parameterization given by $f(s,t)$, which defines the undeformed surface M, and $\tilde{f}(s,t)$, which defines the deformed surface $\tilde{M}$ (Figure 2). Such parameterization can be constructed using standard approaches [SPR06]. f and $\tilde{f}$ may be defined on any type of domain D, e.g., a plane or a sphere, but there must exist a bijection between the surface parameterizations and there must exist a locally planar parameterization. More specifically, let (s,t) be a local parameterization $\alpha : D \to \Re^2$ around $p \in D$, with $\alpha(p) = (s,t)$. Then $f(s,t)$ corresponds to $\tilde{f}(\phi(s,t))$, where $\phi : \Re^2 \to \Re^2$ is a planar bijection over the area of interest. We have the initial correspondence from the undeformed surface M to the deformed surface $\tilde{M}$ (Figure 2):

$$\tilde{M} = \tilde{f}(\phi(f^{-1}(M))) \tag{5}$$

This means, for every point p on M we know the corresponding point $\tilde{p}$ on $\tilde{M}$. ‡

Using the geometric description of strain (Section 3), we directly measure the strain on the surface by mapping a circle of points on M to their corresponding (elliptic) shape on $\tilde{M}$. We take advantage of the fact that strain is zero for a rigid body transformation to map the problem from 3D to 2D. Given a point p on f and a corresponding point $\tilde{p}$ on $\tilde{f}$, we will align the tangent frames § at both p and $\tilde{p}$ with the x,y,z axes in order to reduce the problem to a 2D one (Fig. 4). To achieve this, we first align the tangent frame at p with that at $\tilde{p}$. Then, translate both points to the origin. Next, find the rotation that best aligns both surface normals with the z axis and the s tangent direction with the x axis. We then map the surface to the x,y plane by preserving geodesics as best as possible. At this point, we can drop the z dimension and treat this as a 2D problem (Figure 4).

Let q_i be a set of sample points on M found by projecting a circle around p onto the surface. Map the q_i and p to $\tilde{M}$ to create $\tilde{q}_i$ and $\tilde{p}$ (Fig 4a). Define R_a as a rotation that aligns the tangent frame at point p to that at point $\tilde{p}$ (Fig 4b). We solve for R_a by finding the rotation that best aligns q_i with $\tilde{q}_i$ [Ume91].

$$A = [q_i]^T [\tilde{q}_i] = UDV^T \tag{6}$$

$$R_a = UV^T \tag{7}$$

where A is a 3×3 matrix.

The matrix A that maps q_i to $\tilde{q}_i$ can be expressed as $A = F_{2D}R_a$. To solve for the actual strain tensor F_{2D} we first rotate both sets of points so they lie in the x,y plane as best as possible. Define R_{xyz} to be this rotation for the points $\tilde{q}_i$ (Fig 4c). We can take the point q_i to the x,y plane by first applying R_a, and then R_{xyz}. Now we need to "lift" the points q_i and $\tilde{q}_i$ onto the plane to reduce the problem to 2D. Since we want to measure the change in the geodesics, we approximate the geodesics by placing the points on the plane so that they have both the same proportional angle and the same length as their 3D counterparts. We now solve for the symmetric 2D matrix F_{2D} that best takes the 2D points q_i to the 2D points $\tilde{q}_i$ in a least-squares sense:

$$\begin{bmatrix} (\tilde{q}_i)_x & (\tilde{q}_i)_y & 0 \\ 0 & (\tilde{q}_i)_x & (\tilde{q}_i)_y \end{bmatrix} \begin{bmatrix} F_{1,1} \\ F_{1,2} \\ F_{2,2} \end{bmatrix} = \begin{bmatrix} (q_i)_x \\ (q_i)_y \end{bmatrix} \tag{8}$$

Because F_{2D} is symmetric, the SVD yields $F_{2D} = R_\theta^T S R_\theta$.

This approach solves for the 3D alignment (R_a) indepen-

‡ In Section 5, we will modify ϕ to effectively modify the correspondence between the two surfaces M and $\tilde{M}$.

§ We define the tangent frame at the point p as the partial derivatives of $f(s,t)$ with regard to s and to t, and the normal at p.

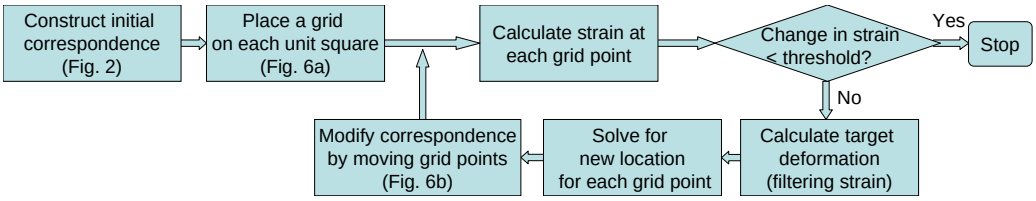

Figure 5: *Procedure for relaxing strain by adjusting the correspondence between between two open regions.*

dently from the deformation (F_{2D}). However, the bias caused by anisotropic scaling of q_i prevents this procedure from reaching an accurate alignment in one single step. Hence, we iterate, scaling q_i by the latest F_{2D} before solving for R_a. In other words, we use:

$$R_a^T R_{xyz}^T F_{2D}^* R_{xyz} R_a q_i \tag{9}$$

(making F_{2D} a 3x3 matrix F_{2D}^*) instead of q_i in Eq. 6. Experimentally, we have found this stabilizes with three iterations, increasing the accuracy by an order of magnitude.

Discussion: We have also experimented with solving for a deformation tensor in three dimensions (F_{3D}), then using the Eigen vectors to determine which two Eigen values correspond to the planar stretch. This is unsatisfactory for two reasons. First, while one of the Eigen vectors usually aligns with the z axis, in the case where one of the planar stretch directions is close to one it becomes unstable and two of the vectors rotate so that they are diagonal. Second, this actually ignores any stretch out of the plane in the z direction.

Our current geodesic approximation is not ideal because it will always be shorter than the real geodesic between p and q_i. A better (but more expensive) solution would be to approximate the true geodesic by, for example, fitting a curve to the two points.

5. Strain relaxation

The previous section explained how we calculate the strain induced by an initial mapping between two surfaces. The resulting strain is, however, not only caused by the deformation itself, but also due to the mapping. As explained by the rubber band analogy, for the region between features, we would like the strain to be spread out evenly instead off accumulated at certain hotspots. In this section, we describe how we modify the correspondence so that the strain is spread more evenly across the deformed surface. This is equivalent to filtering out the noise in the strain across the surface, which means the deformation at a point on the surface should be the averaged deformation of its neighbourhood.

We consider the problem of relaxing the strain between two regions of the surfaces with open boundaries. We first map each region to a unit square in the planar domain (Fig 2) using standard approaches [SPR06]. We then lay identical grids on top of both parameterizations. At this time, the mapping ϕ between one unit square and the other is the identity mapping. As stated previously, the mapping between these two regions (Eq. 5) can be adjusted by modifying the second grid, which essentially modifies ϕ (Fig 6). To allow us to overcome local minima, we use a hierarchical grid-based approach: start with a coarse grid over the entire planar region, then repeat it on a finer grid. At each stage we check that the grid has not folded, taking a smaller step if necessary.

The relaxation procedure constitutes the following steps (see Fig. 5). First, we lay a uniform grid on the planar parameterization of the undeformed surface. We then map it to both the deformed and undeformed surfaces using the initial correspondence. Next, we compute the strain induced by this mapping at each grid point (Section 4). The target deformation at each grid point is computed by averaging the strain tensor at this point with those at its neighbors (Section 5.1). We then construct a linear system to solve for the displacement of the grid points such that the target deformation is achieved (Section 5.2). Finally, we move these grid points to their new locations. We then adjust the parameterization $\tilde{f}$ by moving the projected vertices with the grid. This effectively modifies the correspondence. This procedure is repeated until the change in the total strain energy E falls under a predefined threshold.

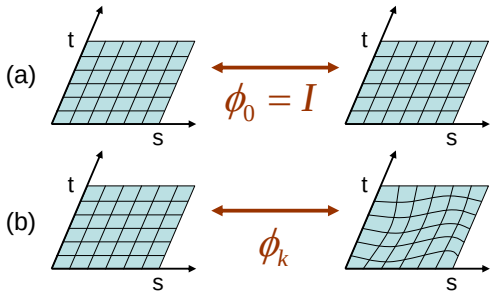

Figure 6: *(a) At the beginning, the initial mapping ϕ_0 is the identity mapping. (b) The mapping ϕ_k as modified after k iterations, effectively changing the mapping between the two surface regions.*

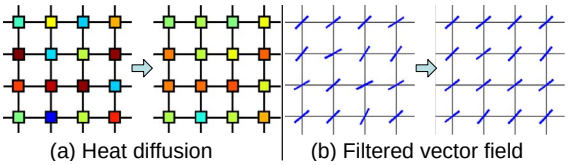

Figure 7: *(a) Magnitude of the deformation can be thought of as heat. The heat at the center point should be the average of that at the surrounding points. (b) We can think of the direction of deformation across the surface as a vector field, with a direction defined at each point on the surface. Then, to average the direction of deformation, the resulted vector field should be a filtered, version of the initial vector field.*

5.1. Filtering the strain tensor

Since the strain tensor captures both the magnitude and the direction of deformation, to filter the strain tensor we average the direction and magnitude of the deformation separately (Fig. 7).

Let u be the current grid point, and v_k its eight neighbor grid points on the first surface, and $\tilde{u}$, $\tilde{v_k}$ be those on the second surface. If F_{2D} is the strain tensor at a grid point, we can decompose this tensor matrix into 3 component matrices as shown in Eq. 1. The average magnitude of deformation $\langle E_i \rangle$ at $\tilde{u}$ is given by:

$$\langle E_i \rangle = \frac{1}{8} \sum_{k=1}^{8} E_{ik} \tag{10}$$

where E_{ik} is the magnitude of deformation at the k^{th} neighbor $\tilde{v}_k$. The average direction of deformation, $\langle \theta \rangle$, is the principal direction computed from Principal Component Analysis (PCA) of all deformation directions at the eight neighbors. Finally, the average deformation, $\langle F_{2D} \rangle$, is the deformation constructed from $\langle E_i \rangle$ and $\langle \theta \rangle$, as shown in Eq. 3 for $i = 1,2$. The resulting strain tensor $\langle F_{2D} \rangle$ is the target strain tensor for the point $\tilde{u}$.

5.2. Constructing the linear system

We would now like to move $\tilde{u}$, $\tilde{v}_i$ by $d\tilde{u}$, $d\tilde{v}_i$ respectively, such that the deformation of the first surface into the second one can be captured by the strain tensor $\langle F_{2D} \rangle$ calculated above. This means that, after both sets of points $u \cup v_i$ and $\tilde{u} \cup \tilde{v}_i$ are translated then rotated such that u and $\tilde{u}$ lie at the origin, and v_i and $\tilde{v}_i$ lie in the (x,y) plane (by applying the rigid transformation matrices $R_1 = R_{xyz} R_a$ and $R_2 = R_{xyz}$ to $u \cup v_i$ and $\tilde{u} \cup \tilde{v}_i$ respectively), $\langle F_{2D} \rangle$ should deform the first set v_i into the second set $\tilde{v}_i$. This transformation is captured in the following equation:

$$R_1(v_i - u) - \langle F_{2D} \rangle R_2[(\tilde{v}_i + d\tilde{v}_i) - (\tilde{u} + d\tilde{u})] = 0 \tag{11}$$

By algebraic transformation, we have:

$$\langle F_{2D} \rangle R_2(d\tilde{v}_i - d\tilde{u}) = R_1(v_i - u) - [\langle F_{2D} \rangle R_2(\tilde{v}_i - \tilde{u})] \tag{12}$$

This is a linear system in the form of $Ax = B$, in which the displacement vectors $d\tilde{u} \cup d\tilde{v}_i$ are the unknowns. Set $LHS = \langle F_{2D} \rangle R_2$ and $RHS = R_1(v_i - u) - [\langle F_{2D} \rangle R_2(\tilde{v}_i - \tilde{u})]$. Now we can construct a sparse linear system to solve for the displacement of the grid points.

$$\begin{bmatrix} 0 & -LHS & \dots & 0 & 0 \\ 0 & 0 & \dots & LHS & 0 \end{bmatrix} \begin{bmatrix} \dots \\ d\tilde{u} \\ \dots \\ d\tilde{v}_i \\ \dots \end{bmatrix} = \begin{bmatrix} RHS \\ RHS \end{bmatrix} \tag{13}$$

For each grid point there are $n+1$ rows of equations (the grid point itself and its n neighbors).

5.3. Constraints

In case we want to restrict movement of certain grid points (due to boundary conditions) or to attract them to particular locations (due to feature correspondence), we can add constraints to the system such as $d\tilde{u}_k = \vec{w}$ (in which $\vec{w} = 0$ for movement restriction). In our experiment, we restrict vertices from moving off of the boundary of the surface.

The constraint for a feature point is constructed as:

$$R_1 u_k - R_2(\tilde{u_k} + d\tilde{u_k}) = \vec{w} \tag{14}$$

Hence:

$$R_2(d\tilde{u_k}) = R_1 u_k - R_2 \tilde{u_k} - \vec{w} \tag{15}$$

which adds one row to the system of equations.

Future work: So far we have explained the strain relaxation procedure for two open regions. To account for arbitrary closed topologies, we can divide the entire surfaces into smaller open regions and perform strain relaxation locally within each region following the same procedure. To prevent discontinuity across the region boundaries, we would redivide the surfaces into different regions overlapping the previous boundaries before relaxing the strain in an iterative process.

Discussion: Using a grid is useful for two reasons. First, it is trivial to implement a coarse to fine hierarchical algorithm. Second, it gets around any issues with poor quality, or greatly differing resolution, meshes. Note that re-meshing in this scenario is not really practical because the target mesh would have to be re-meshed at every time step as the correspondence changes.

6. Experiments

In this section we show the results of applying our approach for surface correspondence using strain relaxation. We start with two sets of synthesized data for a simple cylindrical surface to illustrate the effectiveness of our approach. These are followed by an experiment on real biological data corresponding to cardiac motion in a chick embryo. Our first ex-

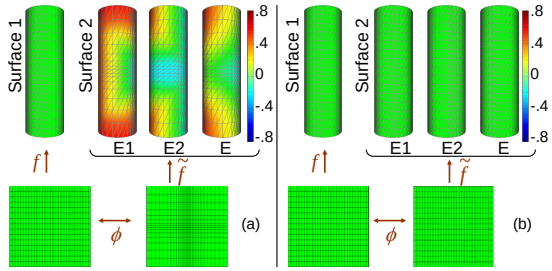

Figure 8: *(a) Strain induced on a cylinder due to incorrect mapping. (b) The final strain after relaxation.*

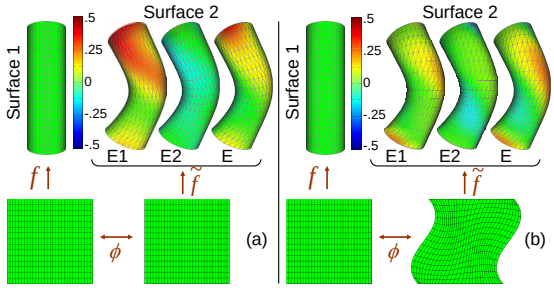

Figure 9: *(a) Strain induced on a cylinder after it has been bent then twisted. (b) The final strain after relaxation.*

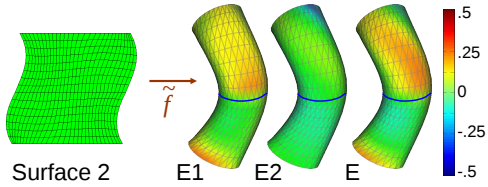

Figure 10: *The final strain after relaxation with constraints: the ring (blue curve) near the middle of the deformed cylinder restricts the middle from moving.*

periment (Fig. 8) is the trivial case of an undeformed cylinder mapped to itself. In this case, we know that the mapping between the two surfaces should be the identity, ie, the two grids should be indistinguishable. We artificially skew the grid on the second planar parameterization to mimic an undesirable (incorrect) mapping between the surfaces (see lower half of Fig. 8a). Fig. 8 shows the strain induced on the second surface as a result of the mapping (a) before and (b) after strain relaxation. With the initial mapping, there is considerable non-uniformity in the strain on the second surface. This is entirely a result of the skewed mapping, since there is no deformation of the surface. As expected, after strain relaxation the mapping between the two grids is the identity (within numerical error). Fig. 8b shows the resulting uniform (and near-zero) strain energy on the second surface.

Our second experiment is on a bent and twisted cylinder. Unlike the previous case, the deformed cylinder is expected to show some strain induced as a result of the deformation. Fig. 9a shows the initial strain on the surface. To perform strain relaxation in this case, the points on the circumference at the two ends of the cylinder are constrained so that they can only slide along the boundaries at the respective ends. After strain relaxation it is observed that the variation of strain energy on the deformed surface is spread out more evenly (Fig. 9b). The strain is now concentrated where the cylinder is bent, with higher strain on the outer section of the bend and lower strain on the inner section. The lower half of Fig. 9b shows the modified mapping between the two surfaces as a result of the strain relaxation procedure. Finally, we show the effect of features on the surface by adding a synthetic constraint to the deformed cylinder. In this case, we pin down the points that lie on the circle at the midpoint of the cylinder. Hence, these points cannot move during strain relaxation ($\vec{w} = 0$). Fig. 10 shows the resulting strain on the deformed surface, as well as the final modified mapping between the two surfaces, as a result of strain relaxation.

Performance: The strain relaxation procedure on this cylinder surface with 480 vertices took about 3 minutes on an Intel Core 6600 (2.4 GHz) processor with 2 GB RAM (MATLAB implementation).

Our last experiment is related to the study of cardiac morphology and function during development, which requires correspondence between data captured at different time points. Our experiments are conducted on image data of the cardiac outflow tract (OFT) of a chick embryo (cardiac development in the chicken is similar to that in humans [LWTR09]) captured *in vivo* using optical coherence tomography (OCT) technique. The OFT is the distal region of the embryonic heart connecting the ventricle with the arterial system. It functions as a primitive valve by contracting to limit blood flow regurgitation. A large portion of congenital heart defects originate in the OFT. The data set reconstructed from the original OCT data comprises three surfaces - the external myocardium surface, internal myocardium surface and the lumen-wall interface. Each surface was captured at three different time points during the cardiac cycle - when the OFT is fully closed, opening, and fully opened (Fig. 11). A quadrilateral mesh of the same topology was constructed from contours for each surface, which gives us the initial correspondence between these data sets. We calculate, and then relax, the strain induced by this mapping on different sets of temporal data for the OFT surfaces.

Figure 12 shows the effect of strain relaxation on the mapping between the lumen-wall interfaces when the OFT is closed (left column) and when it is opening (right column). The strain induced by the initial mapping between these surfaces is shown in Figure 12a. The variation in color shows that the strain is not distributed evenly on the surfaces. In Figure 12b, we show the strain after the mapping has been

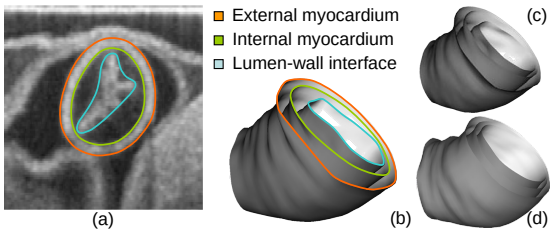

Figure 11: *Surfaces reconstructed from OCT data: Original OCT data with contours (a). Surfaces of the OFT when it is fully closed (b), opening (c), and fully opened (d).*

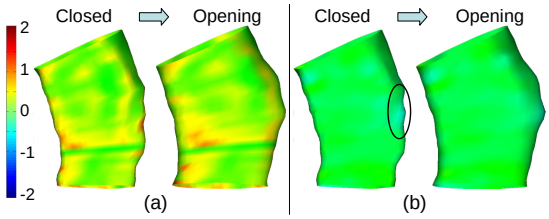

Figure 12: *Strain relaxation on the lumen-wall interface of the outflow tract. The color denotes the magnitude of strain. (a) Artificial strain (hot spots) induced by the initial mapping. (b) The hot spots seen in Fig. 12a have been relaxed. The higher strain (blue region) is expected due to actual deformation of the surface captured.*

adjusted according to the procedure described in Section 5. All the artificial strain hot spots evident in the initial mapping have been evened out. The higher strain that appears on the side is expected due to the actual expansion of the surface captured.

Performance: The strain relaxation procedure on the heart surface with 1000 vertices took about 14 minutes on an Intel Core 6600 (2.4 GHz) processor with 2 GB RAM.

7. Conclusions

In this paper we have shown how to adapt mechanical strain for the purposes of constructing a shape correspondence that has physical meaning, ie, a correspondence that minimizes the change in strain. We present a straightforward technique for calculating strain for a 2D surface embedded in 3D when the source and destination shapes do not have the same mesh topology. This technique makes it possible to phrase the strain minimization algorithm as a sparse matrix least-squares solve. The sparse matrix approach is easily adapted to include feature constraints. Our results show that this approach results in a physically meaningful correspondence between surfaces even in areas lacking identifiable features. We demonstrate the utility of this approach in tracking the motion of biological tissue and internal organs.

Acknowledgements: Funded in part by NSF grants CCF 0702662, DBI 1053171 and DBI 1052688.

References

[ACP03] ALLEN B., CURLESS B., POPOVIC Z.: The space of human body shapes: reconstruction and parameterization from range scans. *ACM SIGGRAPH* (2003), 587–594. 2

[BMTY05] BEG F., MILLER M., TROUVE A., YOUNES L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision 61*, 2 (2005). 2

[EDD*01] ESSEN D. C. V., DRURY H. A., DICKSON J., HARWELL J., HANLON D., ANDERSON C. H.: An integrated software suite for surface-based analyses of cerebral cortex. *Journal of the American Medical Informatics Association 8*, 5 (2001), 443–459. 2

[FKBT08] FILAS B. A., KNUTSEN A. K., BAYLY P. V., TABER L. A.: A new method for measuring deformation of folding surfaces during morphogenesis. *Journal of Biomechanical Engineering 130*, 6 (2008), 061010. 2

[GHDS03] GRINSPUN E., HIRANI A. N., DESBRUN M., SCHRODER P.: Discrete shells. *SCA '03 Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2003), 62–67. 2

[GSS99] GUSKOV I., SWELDENS W., SCHRÖDER P.: Multiresolution signal processing for meshes. *ACM SIGGRAPH* (1999), 325–334. 2

[HPS08] HORMANN K., POLTHIER K., SHEFFER A.: Mesh parameterization: theory and practice. In *ACM SIGGRAPH ASIA 2008 courses* (2008), SIGGRAPH Asia '08, pp. 47:1–47:87. 2

[KCG*10] KNUTSEN A. K., CHANG Y., GRIMM C. M., PHAN L., TABER L. A., BAYLY P. V.: A new method to measure cortical growth in the developing brain. *Journal of Biomechanical Engineering 132*, 10 (2010), 101004. 2

[KS04] KRAEVOY V., SHEFFER A.: Cross-parameterization and compatible remeshing of 3d models. *ACM Transactions on Graphics 23* (August 2004), 861–869. 1

[LWTR09] LIU A., WANG R., THORNBURG K. L., RUGONYI S.: Efficient postacquisition synchronization of 4-d nongated cardiac images obtained from optical coherence tomography: application to 4-d reconstruction of the chick embryonic heart. *Journal of Biomedical Optics 14*, 4 (2009), 044020. 7

[SPR06] SHEFFER A., PRAUN E., ROSE K.: Mesh parameterization methods and their applications. *Foundations and Trends in Computer Graphics and Vision 2*, 2 (2006), 105–171. 1, 2, 4, 5

[STD*07] SHI Y., THOMPSON P. M., DINOV I., OSHER S., TOGA A. W.: Direct cortical mapping via solving partidal differential equations on implicit surfaces. *Medical Image Analysis 11*, 3 (2007), 207–223. 2

[Tab04] TABER L. A.: *Nonlinear Theory of Elasticity: Applications in Biomechanics*. World Scientific Publishing Company, Singapore, 2004. 3

[Ume91] UMEYAMA S.: Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. Pattern Analysis and Machine Intelligence 13*, 4 (1991), 376–380. 4

[WGC*08] WANG Y., GU X., CHAN T. F., THOMPSON P. M., YAU S.-T.: Brain surface conformal parameterization with the slit mapping. *ISBI'08 International Symposium on Biomedical Imaging: From Nano to Micro* (2008), 448–451. 2