

Automatic Alignment of Genus-Zero Surfaces

Patrice Koehl and Joel Hass

Abstract—A new algorithm is presented that provides a constructive way to conformally warp a triangular mesh of genus zero to a destination surface with minimal metric deformation, as well as a means to compute automatically a measure of the geometric difference between two surfaces of genus zero. The algorithm takes as input a pair of surfaces that are topological 2-spheres, each surface given by a distinct triangulation. The algorithm then constructs a map f between the two surfaces. First, each of the two triangular meshes is mapped to the unit sphere using a discrete conformal mapping algorithm. The two mappings are then composed with a Möbius transformation to generate the function f . The Möbius transformation is chosen by minimizing an energy that measures the distance of f from an isometry. We illustrate our approach using several “real life” data sets. We show first that the algorithm allows for accurate, automatic, and landmark-free nonrigid registration of brain surfaces. We then validate our approach by comparing shapes of proteins. We provide numerical experiments to demonstrate that the distances computed with our algorithm between low-resolution, surface-based representations of proteins are highly correlated with the corresponding distances computed between high-resolution, atomistic models for the same proteins.

Index Terms—Conformal mapping, mesh warping, Möbius transformation, nonrigid registration



1 INTRODUCTION

FINDING efficient algorithms to describe, measure, and compare shapes is a central problem in image processing. This problem arises in numerous disciplines that generate extensive quantitative and visual information. Among these, biology occupies a central place. For example, registration of brain anatomy is essential to many studies in neurobiology [1], [2], [3]. In parallel, the belief in molecular biology that the structure (or shape) of a protein is a major determinant of its function has led to the development of many methods for representing, measuring, and comparing the shapes of protein structures [4], [5], [6], [7].

In general, methods that compare shapes can be classified into two categories: those that derive features (also called shape descriptors) for each shape separately that can then be compared using standard distance functions, and those that directly attempt to map one shape onto the other, thereby providing both local and nonlocal elements for comparison. In this paper, we are concerned with the latter. More specifically, we restrict ourselves to methods that generate mappings between two shapes that are defined by surfaces of genus zero.

Surface-based shape comparison techniques aim at defining directly a map between any two surfaces that is as close to an isometry as possible. There have been many methods developed to find such mappings. These methods usually rest on 1) the definition of a distance measure that evaluates how close the map is to an isometry, 2) choices of

sets of landmark points on the two shapes, and 3) an algorithm for finding the mapping between these sets of points that minimizes the distance measure. Note that item 2) is optional, as described below.

An isometry, or map that precisely aligns two surfaces with no distortion, preserves both angles and distances. When the shapes are different, then no isometry can be found, and so some metric distortion is necessary in any alignment. The harmonic or Dirichlet energy [8], [9], the Procrustes distance and its continuous variants [10], the Gromov-Hausdorff distance and variants [11], [12], and the conformal Wasserstein distance and variants that mimic mass transportation [13], [14], [15] are popular metrics used in this context.

Ideally, landmark points should be homologous between the two shapes, should conserve their relative positions, should provide adequate coverage of the shape, and should be found reliably and consistently [16]. In the case of the human cortex for example, landmark points are usually chosen to follow its sulci and gyri patterns (in some cases, the whole curves are considered). The task of finding such meaningful landmark points is most successfully performed manually by skilled technicians with many years of training, with the resulting danger of variability between operators. Many methods have therefore been developed, however, to either automate this process or to circumvent the need to use specific point correspondence in the registration procedure. Methods that fall in the former category rely either on geometric properties of the surface such as critical points identified in the process of flattening the surface [15], or on existing knowledge, such as an atlas for the shape of interest [17]. It is worth mentioning that these methods work on the premise that knowledge of a mapping on a small number of correspondences can be extended to give the full map between the two surfaces of interest [11], [18], [19]. Landmark-free methods, on the other hand, rely on a geometric representation of the surface in which each vertex is assigned a signature, under the

- P. Koehl is with the Department of Computer Science and Genome Center, University of California, Davis, CA 95616. E-mail: koehl@cs.ucdavis.edu.
- J. Hass is with the Department of Mathematics, University of California, Davis, CA 95616. E-mail: hass@math.ucdavis.edu.

Manuscript received 22 Aug. 2012; revised 27 Jan. 2013; accepted 20 July 2013; published online 29 July 2013.

Recommended for acceptance by P. Golland.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2012-08-0652.

Digital Object Identifier no. 10.1109/TPAMI.2013.139.

premise that points with similar signatures are more likely to correspond. Spectral techniques fit in this category. Rustamov, for example, introduces the global point signature (GPS) of a point on the shape, which encodes both the eigenvalues and the eigenfunctions of the Laplace-Beltrami operator evaluated at that point [20]. Guibas et al. introduced the heat kernel signature, a similar robust and multiscale invariant defined on each point of the surface. This invariant is found by solving a partial differential equation involving the same Laplace-Beltrami operator [21]. Fischl et al. assigned to each vertex on a mesh an average convexity, which is then used to drive the alignment of cortical surfaces [22], [23]. Vaillant and Glaunès introduced a representation of surface in the form of currents and then imposed a Hilbert space structure on it, whose norm is used to quantify the similarity between two surfaces [24].

In the special case that the two surfaces to be mapped are of genus zero (i.e., the surfaces are topologically equivalent to the 2-sphere), two alternate approaches have been used to generate the closest to-isometric mapping, namely, those that apply first a parameterization (or mapping) of the surfaces onto the sphere, and those that directly align the surfaces in their own coordinates (usually the 3D Cartesian coordinates).

The existence of a conformal mapping of a genus-zero surface to the round sphere is guaranteed by the Uniformization Theorem [25], [26]. Various methods have been proposed to generate such a conformal mapping in the discrete case (i.e., when the surface is represented by a mesh), including building a linear system that approximates the Laplace-Beltrami operator [27], using circle packing [28], solving for degree one harmonic maps [29], [30], or minimizing an angle-based functional [31]. Once two genus-zero surfaces have been mapped conformally onto the sphere, the search for (near) isometries between them can be made more tractable by restricting to a search within the Möbius group, the group of bijective conformal self-mappings of the sphere. The Möbius group acts on the sphere with six degrees of freedom. A single Möbius transformation, and therefore a conformal mapping between the two surfaces, is determined by specifying correspondences between exactly three points. Lipman and Funkhouser implemented this idea by sampling random triplets of points over each surface, computing the Möbius transformation defined by those triplets, and by “voting” over the samples, using as a ranking criteria the estimated deviation from isometry [32]. This idea was further extended to automatically quantify the overall similarity between surfaces [15]. It was also previously used by Tosun et al. [33] and Gu et al. [30] to minimize sulcal distances when aligning brains.

A conformal mapping between two surfaces of genus zero can only align exactly three points; exact matching of more than three landmark points therefore requires relaxation of the constraints of angle preservation. Wang et al. [34], for example, introduced an approach for comparing two genus-zero surfaces C_1, C_2 that balanced conformability with landmark correspondence. They computed a conformal map $f_1 : C_1 \rightarrow S^2$ and then searched for a second map $f_2 : C_2 \rightarrow S^2$ that minimizes an energy function with two

terms $E_1 + E_2$. The first term E_1 is the standard Dirichlet energy, whose minimizer gives rise to a conformal map. The second term is given by $E_2 = \sum \|f_2(q_i) - f_1(p_i)\|^2$, where the norm represents the euclidean distance of two points on the round sphere, and $\{p_i \in C_1\}$ and $\{q_i \in C_2\}$ are matching landmark points. Joshi et al. introduced a similar extended cost function that includes a strain energy (based on the Laplace-Beltrami operator) and a penalty term for sulci matching, with the alignment being performed on a flat surface instead of the sphere [35].

In this work, we are interested in generating a *globally optimal conformal mapping* between two surfaces of genus zero. In contrast to the works of Wang et al. [34] and Joshi et al. [35] mentioned above, we do not relax the constraints of angle preservation; instead, we propose a new algorithm that fully eliminates the use of landmarks, thereby generating a mapping based only on the knowledge of the two surface triangulations. In this approach, the whole mesh representing the source surface is warped onto the target surface, using the mapping defined through the composition of discrete conformal mappings of the surfaces onto the sphere and the Möbius transformation between these mappings. The discrete mappings onto the sphere are generated using the algorithm introduced by Springborn et al. [31]. The Möbius transformation is then optimized by gradient descent to lead to minimal distortion between the source mesh and its image, where distortion is measured as difference from isometry. A good choice of initial maps is essential for success for this kind of method, as one expects many local minima to exist in the 6D space of conformal maps that we explore. We generate initial alignments based on choices of triplets of points in each surface computed from a best ellipsoid approximation for each. The success of such an approach is likely to be problem dependent. We note that our optimal conformal mapping is a composition of Möbius transformations; as each Möbius transformation is a homeomorphism, it is therefore a homeomorphism between the two surfaces.

This paper is organized as follows: Section 2 provides the mathematical background for our algorithm: discrete conformal geometry and distance measures between meshes. In Section 3, we provide the details of its implementation, focusing on how to ensure its robustness. Section 4 presents the results obtained by our algorithm on automatic registrations of brains, as well as on comparing surfaces of protein structures. We conclude the paper with a discussion on future developments.

2 BACKGROUND

2.1 Basic Ideas

Let S_1 and S_2 be two surfaces of genus zero, represented by the meshes \mathcal{M}_1 and \mathcal{M}_2 , respectively. Both meshes are taken to be triangular, with $\mathcal{M} = (V, E, T)$ and $V = \{v_i\}$, $E = \{e_{ij}\}$, and $T = \{t_{ijk}\}$ denoting the vertices, edges, and triangles, respectively. We note that these two meshes are independent of each other and are likely to have different combinatorics. Our goal is to define a discrete map $f : \mathcal{M}_1 \rightarrow f(\mathcal{M}_1)$ where $f(\mathcal{M}_1)$ is a geometric realization of \mathcal{M}_1 onto S_2 , that is as close as possible to an isometry, i.e., that




Fig. 1. Globally optimal conformal mapping.

minimizes the distortion of pairwise distances between vertices. The key observation [32] that makes the problem tractable is that f can be understood as the composition of three discrete conformal mapping functions, C_1 , m , and C_2^{-1} (see Fig. 1). The group of conformal self-mappings of the sphere is well understood and is called the Möbius group. Any transformation m in this group is defined by specifying the image of three points, and thus has six degrees of freedom. The mapping f is then constructed by optimizing m so that the composition $C_2^{-1} \circ m \circ C_1$ is as near to an isometry as possible. We use this structure to develop an automatic algorithm for comparing two genus-zero surfaces. This method involves two main steps: 1) generation of parameterizations C_1 and C_2 of the meshes \mathcal{M}_1 and \mathcal{M}_2 onto the sphere, and 2) optimization of the Möbius transformation m . Mathematical details for each of these two steps appear in the next sections.

2.2 Discrete Conformal Mapping to the Sphere

While Riemann's Uniformization Theorem guarantees that any smooth genus-zero surface S can be mapped conformally (with angles preserved) to the unit sphere, in applications we are forced to work with discrete approximations of these underlying smooth objects. The theoretical underpinnings of the theory of discrete conformal maps are still being developed, but many methods have been developed to compute them in practice. We follow the approach proposed by Springborn et al., which introduces a notion of discrete conformal equivalence [31]. While we refer the reader to their paper for a full description of this approach, we summarize it here to introduce definitions and equations relevant to our algorithm. Derivations of the formulas given below can be found in their paper.

Let us consider a triangular mesh $\mathcal{M} = (V, E, T)$ embedded in a 2D manifold. We do not restrict its topology or assume the presence of a boundary. The intrinsic geometry of \mathcal{M} is encoded in its edge lengths.

Definition 1. A discrete metric on M is a function l defined on the set of edges E that assigns to each edge e_{ij} a length l_{ij} such that the triangle inequalities are satisfied for all triangles in T , so that no side of a triangle has length longer than the sum of the other two side lengths.

The notion of discrete conformal equivalence is defined as follows.

Definition 2. Two discrete metrics l and \tilde{l} on M are discretely conformally equivalent if, for some assignment of numbers u_i to the vertices v_i , the metrics are related by the formula:

$$\tilde{l}_{ij} = e^{(u_i + u_j)/2} l_{ij}. \quad (1)$$

Starting with some discrete metric l , we are interested in finding a new discrete metric \tilde{l} that is discretely conformally equivalent and that has particularly nice geometric properties. In the following, we set l to be euclidean distance, meaning that the relations between lengths and angles of triangles are those of the classical euclidean case. Note that other metrics are possible. Given such a metric l and a triangle t_{ijk} in M , the angle α_{jk}^i at vertex v_i opposite edge e_{jk} can be recovered from the lengths of the sides of the triangle by standard euclidean trigonometry, giving the formula:

$$\alpha_{jk}^i = 2 \tan^{-1} \sqrt{\frac{(l_{ij} + l_{jk} - l_{ki})(l_{jk} + l_{ki} - l_{ij})}{(l_{ki} + l_{ij} - l_{jk})(l_{jk} + l_{ki} + l_{ij})}}. \quad (2)$$

The curvature K_i at an interior vertex v_i is given by the excess angle sum:

$$K_i = 2\pi - \sum_{v_i \in t_{ijk}} \alpha_{jk}^i. \quad (3)$$

The problem of mesh parameterization can then be stated as finding weights u_i such that the new corresponding metric \tilde{l} has curvature 0 at all interior vertices in the mesh. These conditions lead to a system of nonlinear equations, one per interior vertex. Springborn et al. [31] have shown that if a solution to this system of equations exists, it can be found as the *unique* minimizer of the convex energy in u :

$$CE(u) = \sum_{t_{ijk} \in T} f(u_i, u_j, u_k) + \pi \sum_{v_i \in V} u_i, \quad (4)$$

where

$$f(u_i, u_j, u_k) = \frac{1}{2} (\tilde{\alpha}_{jk}^i \tilde{\lambda}_{jk} + \tilde{\alpha}_{ki}^j \tilde{\lambda}_{ki} + \tilde{\alpha}_{ij}^k \tilde{\lambda}_{ij}) + \mathcal{J}(\tilde{\alpha}_{jk}^i) + \mathcal{J}(\tilde{\alpha}_{ki}^j) + \mathcal{J}(\tilde{\alpha}_{ij}^k) - \frac{\pi}{2} (u_i + u_j + u_k), \quad (5)$$

with

$$\tilde{\lambda}_{ij} = 2 \log(l_{ij}) + u_i + u_j, \quad (6)$$

and \mathcal{J} is Milnor's Lobachevsky function:

$$\mathcal{J}(x) = - \int_0^x \log |2 \sin t| dt. \quad (7)$$

This can be understood by considering the partial derivatives of the energy:

$$\frac{\partial CE}{\partial u_i} = \frac{1}{2} \left(2\pi - \sum_{t_{ijk} \ni v_i} \tilde{\alpha}_{jk}^i \right), \quad (8)$$

so that $\nabla CE(u) = 0$ if and only if the scaling factors u solve the problem and give a zero curvature metric.

The method described above allows for the flattening of a mesh topologically equivalent to a disk onto a plane.




Fig. 2. *Discrete conformal mapping to the sphere.* The mesh representing the surface S_1 (A) is first mapped to the plane (B) after removal of a vertex v_0 and its open star to make it topologically equivalent to a closed disk. The planar layout is then projected stereographically to the sphere (C), with v_0 being reinstated at the North pole. The spherical mesh is then normalized to ensure zero mass-center (D). Note that the sulci appear clearly on the normalized mesh.

When used in combination with stereographic projection from the plane to the sphere, it can be applied to the problem of mapping meshes with spherical topology onto the sphere with the following algorithm (see also Fig. 2).

Algorithm 1. *Discrete conformal mapping of a genus 0 mesh \mathcal{M} with discrete metric l to the unit sphere [31], [36].*

- (1) Initialization. Select vertex v_0 with smallest curvature. Set $u_j = -2 \log l_{j0}$ for all vertices v_j in the link of v_0 .
- (2) Let \mathcal{M}_0 be \mathcal{M} minus the open star of v_0 . \mathcal{M}_0 is topologically a closed disk.
- (3) Flatten \mathcal{M}_0 by solving iteratively for u_i that minimize the energy $CE(u)$ defined in Section 4.
- (4) Layout \mathcal{M}_0 on the plane under the minimized metric \tilde{l} .
- (5) Project planar layout on the sphere stereographically and reinstates v_0 at the North Pole.
- (6) Apply Möbius normalization to ensure that the center of mass of all vertices is at the origin.

The choice of the vertex v_0 in step (1) is arbitrary; we have found that choosing the vertex with the smallest curvature works best in practice. The minimization in step (3) is performed over all interior vertices of the mesh \mathcal{M}_0 . The vertices at its boundary, i.e., those that correspond to the link of v_0 , have their weights fixed to their values given in step (1).

The result of this procedure, if it exists, is a polyhedron with vertices on the sphere that is discretely conformally equivalent to \mathcal{M} . It may not exist, as the set of nonlinear equations solved in step (3) is not guaranteed to have a solution. Indeed, while the variational energy CE is convex, we do not have a convex optimization problem. The set of weights u_i that results in new edge lengths that satisfy the triangle inequalities is not convex. We can circumvent to some extent this difficulty by extending the domain of CE :

$$\text{if } \tilde{l}_{jk} > \tilde{l}_{ki} + \tilde{l}_{ij} \text{ then } \tilde{\alpha}_{jk}^i = \pi \text{ and } \tilde{\alpha}_{ki}^j = \tilde{\alpha}_{ij}^k = 0.$$

This simple fix still does not guarantee that a solution exists. We will discuss this further in Section 3.

The parameterization resulting from this algorithm is not unique. The choices made in the algorithm can lead to parameterizations that differ by a rotation of the sphere. By imposing the zero mass-center (step (6)), the mapping is at least unique up to the euclidean rotation group.

2.3 Möbius Normalization for Zero Mass-Center

In his elegant paper [37], Springborn proved that for “ $n \geq 3$ distinct points in the d -dimensional unit sphere, there exists a Möbius transformation such that the barycenter of the transformed points is the origin.” He did not, however, include a practical method for building such a transformation. We show here that searching for this transformation can be formulated as a nonlinear optimization problem with four variables.

Let us consider $n \geq 3$ distinct points in the 2D unit sphere. The stereographic projection from the North pole of this sphere onto the xy -plane takes the point v_i with coordinates (x_i, y_i, z_i) to the point:

$$P_i = (X_i, Y_i, 0) = \left(\frac{x_i}{1 - z_i}, \frac{y_i}{1 - z_i}, 0 \right),$$

that can be identified with the complex coordinate $z_i = X_i + iY_i$.

A Möbius transformation m that maps infinity to infinity in the Riemann plane (i.e., that maps the North pole of the sphere to itself) can be represented as a linear form $m(a, b) = az + b$, where a and b are complex numbers and $a \neq 0$. The point P_i is mapped into a point P'_i with complex coordinate $z'_i = az_i + b$. Let us define $\omega'_i = |z'_i|^2 = (az_i + b)(\overline{az_i + b})$. The inverse stereographic projection of P'_i gives the following formula for the coordinates of the image $v'_i = (x'_i, y'_i, z'_i)$ of v_i under the transformation m :

$$v'_i = \left(\frac{az_i + \overline{az_i} + b + \bar{b}}{1 + \omega'_i}, \frac{az_i - \overline{az_i} + b - \bar{b}}{i(1 + \omega'_i)}, \frac{\omega'_i - 1}{\omega'_i + 1} \right).$$

The transformation m satisfies the zero mass center condition if and only if the functional

$$F(a, b) = \frac{1}{n^2} \left(\left(\sum_{i=1}^n x'_i \right)^2 + \left(\sum_{i=1}^n y'_i \right)^2 + \left(\sum_{i=1}^n z'_i \right)^2 \right) \quad (9)$$

satisfies $F(a, b) = 0$. The summations extend over all points, and the coordinates are nonlinear functions of a and b , as given above. We find the best (in the least-squares sense) transformation $m(a^*, b^*)$ by solving

$$(a^*, b^*) = \arg \min_{a, b} F(a, b). \quad (10)$$

2.4 Warping the Source Mesh onto the Destination Surface

Following the idea illustrated in Fig. 1, a mapping f between the source and destination surfaces S_1 and S_2 can be written as the composition $f = C_2^{-1} \circ m \circ C_1$, where C_1 and C_2 are the discrete conformal mappings of S_1 and S_2 on the sphere, respectively, and m is a Möbius transformation. Any such transformation m is characterized by four complex numbers and is given by the closed-form formula:

$$m(z) = \frac{az + b}{cz + d} \quad \text{with} \quad ad - bc = 1. \quad (11)$$

For convenience of notation, we store the four complex numbers a, b, c , and d into a vector \vec{h} and define $m = m(\vec{h})$. Note that \vec{h} contains eight real numbers, but corresponds to six degrees of freedom only, as $ad - bc = 1$.

Let us call \mathcal{M}_1 and \mathcal{M}_2 the genus-zero meshes representing S_1 and S_2 , and \mathcal{SM}_1 and \mathcal{SM}_2 the images on the sphere obtained from the discrete mappings C_1 and C_2 . We generate the warping $W_m(\mathcal{M}_1)$ of \mathcal{M}_1 onto the surface S_2 by locating the successive images v'_i , v''_i , and v'''_i of a vertex v_i in \mathcal{M}_1 for the mappings C_1 , a Möbius transformation m , and the inverse of C_2 , in this order. Both v'_i and v''_i lie on the sphere, with $v''_i = m(v'_i)$. The key idea to finding v'''_i is to locate v''_i in the spherical triangulation \mathcal{SM}_2 and transfer this location onto \mathcal{M}_2 . To speed up this location step, we partition the unit sphere, using uniform subdivisions in the polar coordinates (ϕ, θ) . We then find a representative triangle in \mathcal{SM}_2 for each subdivision. The algorithm for warping \mathcal{M}_1 onto the surface S_2 is given in Algorithm 2.

Algorithm 2. *Warping a genus 0 mesh \mathcal{M}_1 on a destination surface S_2 defined by a mesh \mathcal{M}_2 .*

```

for all vertices  $v_i$  in  $\mathcal{M}_1$  do
  (1) Compute  $v'_i = m(v_i)$ , where  $v'_i$  is the vertex
      equivalent to  $v_i$  in  $\mathcal{SM}_1$ .
  (2) Identify the subdivision of the sphere containing  $v'_i$ 
      and its representative triangle  $t'$ .
  (3) Compute barycentric coordinates  $(\alpha, \beta, \gamma)$  of  $v'_i$  in
      triangle  $t'$ .
  while  $\alpha < 0$  or  $\beta < 0$  or  $\gamma < 0$  do
    (4) Identify vertex  $p$  of  $t'$  corresponding to one of
        the negative barycentric coordinates. Update  $t'$ 
        with its adjacent triangle opposite to  $p$ .
    (5) Update barycentric coordinates  $(\alpha, \beta, \gamma)$ .
  end while
  (5) Compute position of  $v'''_i = f(v_i)$  on the surface  $S_2$ 
      by propagating the barycentric coordinates  $(\alpha, \beta, \gamma)$ 
      onto the triangle  $t$  in  $\mathcal{M}_2$  that corresponds to  $t'$ .
end for

```

The result of this procedure is a new mesh, $W_m(\mathcal{M}_1)$, on S_2 , with the same combinatorics as \mathcal{M}_1 , but different geometry. Note that a different warping $W_m(\mathcal{M}_1)$ is obtained for each Möbius transformation m .

2.5 Measuring the Distance between Two Meshes

The two meshes \mathcal{M}_1 and $W_m(\mathcal{M}_1)$ have the same combinatorics (V, E, T) but different geometries (encoded in the positions of the vertices), as \mathcal{M}_1 is a representation of S_1 while $W_m(\mathcal{M}_1)$ sits on the surface S_2 . By construction, $W_m(\mathcal{M}_1) = f(\mathcal{M}_1)$. We define the *elastic energy* L associated with this mapping f as

$$L(f) = \sum_{e_{ij} \in E} \left(\frac{\|f(v_j) - f(v_i)\|}{\|v_j - v_i\|} - 1 \right)^2, \quad (12)$$

where $\|\cdot\|$ is the usual L_2 norm.

$L(f)$ is a measure of the differences between the two shapes represented by the surfaces S_1 and S_2 .

A mapping f with $L(f) = 0$ is called an “isometry” and the two meshes \mathcal{M}_1 and $W_m(\mathcal{M}_1)$ are then said to be

isometric. Note that this is a weak concept of isometry. Formally, we can only state that if f is an isometry of the underlying surfaces, then $L(f) = 0$.

We work in the class of conformal maps, so angles are preserved. The remaining distortion of a map from an isometry is exactly reflected in the conformal stretching factor. The elastic energy $L(f)$ defined above measures the average stretching of the edges in a mesh and thus reflects the distance of a conformal map from an isometry.

2.6 A General Algorithm for Mapping Two Surfaces of Genus Zero

As illustrated in Fig. 1, we rely on the idea that a conformal mapping f between two surfaces S_1 and S_2 of genus zero can be written as the composition of two discrete conformal mappings C_1 and C_2 that parameterize S_1 and S_2 onto the sphere, and a Möbius transformation m . In optimizing the map produced from this composition, C_1 and C_2 are fixed, while m is variable and depends on six degrees of freedom summarized in a parameter vector \vec{h} (see (11)). The key to our approach is to choose the transformation m to yield the minimum weighted distance between the mesh \mathcal{M}_1 representing S_1 and its image $W_m(\mathcal{M}_1)$ warped onto S_2 . This approach eliminates the need to define correspondences between landmark points on the two surfaces. The weighted distance between a mesh and its image under the conformal mapping f is measured by the elastic energy $L(f)$ of f . For convenience of notation, it is written $L(f) = L(m(\vec{h})) = L(\vec{h})$ as only m is variable. We have developed all the tools we need to perform this optimization, namely: 1) an algorithm for computing the discrete conformal mappings C_1 and C_2 (see Algorithm 1), 2) an algorithm for generating the warping of a discrete mesh onto a surface for a given Möbius transformation m (see Algorithm 2), and 3) a definition of the elastic energy $L(f)$ that measures its distance to an isometry (see (12)).

Simple calculations provide the analytical expressions for the elastic energy function $L(\vec{h})$ and its gradient with respect to \vec{h} . This allows us to apply a steepest descent algorithm to optimize the Möbius transformation m . Our general algorithm for comparing the two surfaces S_1 and S_2 is then:

Algorithm 3. *Correspondence-free comparison of two discrete surfaces of genus zero.*

```

Initialization. Apply algorithm 1 to map  $\mathcal{M}_1$  and  $\mathcal{M}_2$ 
onto the sphere.
(1) Initialize Möbius transformation  $m_0 = m(\vec{h}_0)$ .
for  $n = 0, \dots$  until convergence do
  (2) Generate  $W_m(\mathcal{M}_1)$  using algorithm 2 where
       $m = m(\vec{h}_n)$ .
  (3) Compute  $L(\vec{h}_n)$  and its gradient  $\nabla L(\vec{h}_n)$  with
      respect to  $\vec{h}_n$ .
  (4) Update  $\vec{h}_{n+1} = \vec{h}_n - \alpha_n \nabla L(\vec{h}_n)$ .
  (5) Check for convergence: if  $L(\vec{h}_{n+1}) < \text{TOL}$ , stop.
end for

```

The damping parameter α_n in step (4) is obtained by solving the equation $L(\vec{h}_n + \alpha_n \nabla L(\vec{h}_n)) \leq L(\vec{h}_n)$ using a line search method. The value of TOL is set to a small constant related to machine error.

The result of this procedure is a warping of the mesh \mathcal{M}_1 onto the surface S_2 that minimizes distance from an

isometry among nearby conformal maps, with distance measured by the elastic energy. In addition, it gives a measure of the distance between \mathcal{M}_1 and its warped image that reflects the geometric differences between the two surfaces. In many cases where the surfaces S_1 and S_2 are isometric, the procedure will produce the isometry.

3 IMPLEMENTATION

There are several issues to address in generating a robust implementation of the algorithms discussed above that is fast enough to be usable in practice. In the following, we describe with details the implementations of the main steps of the method. We do not provide a convergence or complexity analysis, and leave those for future work.

3.1 Discrete Conformal Mapping on the Sphere

Flattening a mesh topologically equivalent to a disk. Following the formalism developed by Springborn et al., we have shown in the previous section that the discrete mesh parameterization problem can be recast into an unconstrained convex optimization problem with explicit formulae for the target function $CE(u)$ (see Section 4), its gradient, and its Hessian (see [31]). As noted by Springborn et al., the Hessian is positive semidefinite with only the constant vector in its null-space. To find the vector u that minimizes the extended energy, we use a trust region Newton method as implemented in the program TRON [38].

TRON is globally convergent and will find a vector u that minimizes CE . It does not guarantee, however, that the solution is feasible: It is possible that the minimum is reached for u that defines a new discrete metric \tilde{l} that does not satisfy the triangle inequalities. In all the cases we have tested, we observed that the violations of triangle inequalities, if any, occurred in sliver triangles, for which one of the interior angles is smaller than 5 degrees. This lead us to the following strategy for eliminating the violations:

1. solve the convex optimization problem for u ,
2. compute all edge lengths based on the new metric \tilde{l} ,
3. detect all triangles whose edge lengths violate the triangle inequality and flip the edges opposite to their smaller angle, and
4. repeat points 1-3 until all triangle inequalities are satisfied.

The same strategy of edge flipping was already considered by Springborn et al. [31]. For most of the test cases described in the result section, a single cycle was found to be sufficient; in the remaining cases, two cycles removed all violations.

We note that this simple strategy will not work if the mesh contains regions with high densities of sliver triangles. In such cases, remeshing is required.

Laying out the planar mesh. Once the weights u have been found and it is verified that the corresponding discrete metric \tilde{l} is free of triangle inequality violations, we have new edge lengths that guarantee that the mesh \mathcal{M} is flat, i.e., it can be laid out in the plane after removal of a vertex. We assume that the mesh is oriented, and that the vertex lists for all triangles are ordered with consistent orientation. We generate the coordinates of all N vertices in \mathcal{M} under the

metric \tilde{l} using the following layout procedure. Each triangle in \mathcal{M} is assigned a flag, vt , initially set to zero; a similar flag vv is assigned to each vertex v . We create an empty master list L . We pick a triangle t_0 at random in the mesh structure. One of its vertices i is set at the origin of the plane, and a second one j is placed on the x -axis, with its distance to the origin set to \tilde{l}_{ij} . The triangle t_0 is added to the list L and its flag is set to one; the flags vv of the two vertices i and j are set to one. The layout algorithm proceeds as follows.

Algorithm 4. *Planar layout of the parameterized mesh.*

```

Set  $nv = 0$ 
while  $|L| \neq 0$  do
   $t = \text{pop}(L)$ .
  (1) Let  $t = (i, j, k)$  and  $S(t) = vv(i) + vv(j) + vv(k)$ .
  if  $S(t) < 3$  then
    (2) Let  $i$  be the vertex with  $vv(i) = 0$ . Build  $i$  using
        standard geometry;  $vv(i) = 1$ ;  $nv := nv + 1$ .
  end if
  for all  $t_n$  adjacent to  $t$  with  $vt(t_n) = 0$  do
     $vt(t_n) = 1$ ;  $L = L \cup \{t_n\}$ .
  end for
  if  $nv \% 10,000 = 0$  or  $nv = N$  then
    (3) Regularize layout
  end if
end while

```

In step (1), the triangle t is either t_0 or a neighbor of a triangle that was already laid out; in both cases, at least two of its vertices have already been considered and $S(t) \geq 2$. The construction process in step (2) is then very simple: The positions of two vertices of t in the plane are known, and the distances from these vertices to the third vertex are given by the new edge lengths of the mesh. This is a simple geometric problem that has two solutions, one for each orientation. When building the first triangle, one solution is picked at random. For all other triangles, only one solution is feasible as the orientation of t must match the orientations of its neighboring triangles. In theory, this breadth-first approach will lay out meshes of any size. We have observed in practice, however, that accumulation of numerical errors can lead to significant distortions, especially for meshes with large length ratios. To circumvent this problem, we have introduced a regularization process as step (3). At a stage in the process with nt triangles already built, we can compute a distortion error as follows:

$$DE = \sum_{e_{ij}} ((x_i - x_j)^2 + (y_i - y_j)^2 - \tilde{l}_{ij}^2)^2, \quad (13)$$

where the summation extends over all edges e_{ij} in the nt triangles, and (x_i, y_i) are the planar coordinates of vertex v_i . If DE is small (we use a cutoff of 10^{-8}), the layout is considered correct. Otherwise, we perform a nonlinear optimization of the coordinates of the vertices to reach a minimum of DE . Simple calculations provide explicit formulae for the gradient and Hessian of DE , and we can use the Newton trust region method as implemented in TRON to perform this optimization. To save computing time, regularization is only performed at multiples of 10,000 vertices and when the layout is complete. In all

examples described in the results section, this procedure yields worst relative length errors bounded by 10^{-7} .

Möbius normalization for zero mass-center. As described in Section 2, the normalization is cast into a nonlinear optimization problem (see (10)). An explicit formula for the dependence of $F(a, b)$ on the parameters a and b of the Möbius normalization is available (see (9)). It is also easy to compute its gradient and Hessian. We can therefore again use the Newton trust region method as implemented in TRON to perform this optimization. Typical examples indicate that we only need up to 10 iterations with TRON to reach $\|\nabla F(a, b)\|_2 \leq 10^{-12}$.

3.2 Warping the Source Mesh onto the Target Surface

The method we propose for warping a mesh \mathcal{M}_1 onto a surface S_2 is quite simple. Starting with a vertex v_i in \mathcal{M}_1 , we identify its corresponding point v'_i in the spherical mesh \mathcal{SM}_1 . We then locate its image $v''_i = m(v'_i)$ on the spherical mesh \mathcal{SM}_2 and transfer this location onto the mesh \mathcal{M}_2 . The implementations of these steps is straightforward, with the exception of the point location problem. As this procedure needs to be repeated for all vertices in \mathcal{M}_1 (which can be on the order of hundreds of thousands), and subsequently for many trials for the Möbius transformation m , we need it to be fast and reliable.

Our approach is akin to the “jump-and-walk” algorithm used for point location in 2D or 3D triangulations. Given a point v''_i on the sphere, we first “jump” to a trial triangle t in the spherical mesh \mathcal{SM}_2 . If the (spherical) barycentric coordinates of v''_i with respect to t are all positive, v''_i is inside t and the algorithm stops. If one of the coordinates is negative, we identify the corresponding vertex in t and “walk” toward the triangle adjacent to t that is opposite this vertex. The procedure is then iterated until the correct triangle is identified. It is easy to see that this method is guaranteed to converge. It may not be fast, however, as its speed depends on the quality of the jump, i.e., on how close the initial trial triangle is from the actual triangle containing v''_i . Standard jump-and-walk techniques pick the initial triangle from a random subset. We propose an alternative approach that requires preprocessing. A similar method was proposed by Wu et al. [39].

We first put a grid G on the sphere based on spherical coordinates. We chose 400 divisions for the azimuthal angle ϕ and 200 divisions for the polar angle θ , corresponding to a total of 80,000 rectangular subdivisions on the sphere. We numerically define the coverage of a triangle t in \mathcal{SM}_2 by sampling uniformly 20 points inside t and computing their polar coordinates. The subdivisions of the grid G that contain these points are assigned t as a representative. This procedure is repeated for all triangles in \mathcal{SM}_2 . Any rectangular subdivision s that does not have a representative at this stage is processed further. First, we find the closest subdivision to s with a representative, t_0 , using a breadth-first algorithm. The triangle containing the center of s is then located with the jump-and-walk procedure described above, using t_0 as the initial trial. This triangle is set as the representative of s . As a result of the this procedure, all subdivisions in the Gaussian grid are assigned a representative triangle in \mathcal{SM}_2 . The initial trial

triangle for a vertex v''_i is then taken to be the representative triangle of the subdivision that contains v''_i .

3.3 Optimal Möbius Transformation for Near-Isometry

The Möbius transformation that leads to a closest to isometric mapping among conformal maps between the two surfaces of interest is obtained as the solution of a nonlinear optimization problem (see Algorithm 3). We have adapted a steepest descent approach to solve this problem. Steepest descent methods are usually fast and easy to set up. They are, however, very sensitive to local minima and highly dependent on the quality of the initial guess for the solution.

A random or a trivial initial guess (such as the identity transformation) is likely to lead to a local minimum. We have therefore developed a simple procedure to automatically generate better starting points. The idea is to use the best ellipsoid approximation to each surface to give the initial alignment. While this alignment is specified by setting the images of three points, we do not rely on user selected landmarks or on local geometric features to select these points. Instead, we simply compute the principal components of the sets of points representing the two shapes. For the mesh \mathcal{M}_1 , we compute the covariance matrix over its set of vertices:

$$C_1 = \sum_{v_i} (\vec{v}_i - \vec{\mu})(\vec{v}_i - \vec{\mu})^T, \quad (14)$$

where \vec{v}_i is the vector of coordinates for vertex v_i , μ is the center of mass of the vertices, and the summation runs over all vertices. The ellipsoid whose axes are defined by the unit eigenvectors of C_1 , scaled by the associated eigenvalues, is the best fit ellipsoid to the mesh \mathcal{M}_1 . These three axes cut the surface represented by \mathcal{M}_1 in three pairs of points, (A_1, A'_1) , (B_1, B'_1) , and (C_1, C'_1) . Using the same procedure on \mathcal{M}_2 , we get three corresponding pairs of points, (A_2, A'_2) , (B_2, B'_2) , and (C_2, C'_2) . The three points (A_2, B_2, C_2) defines a direct orientation for \mathcal{M}_2 . There are four choices of triplets of points on \mathcal{M}_1 whose correspondences to these three points lead to an alignment of the axes of the two ellipsoids with proper orientation; these are (A_1, B_1, C_1) , (A_1, B'_1, C'_1) , (A'_1, B_1, C_1) , and (A'_1, B'_1, C'_1) . Each of these correspondences defines a unique Möbius transformation. One of the advantages of choosing these points is that they are well separated, which leads to stability in the corresponding transformation. Each of these transformations is then used as an initial guess for the steepest descent approach. This leads to four different optimizations. The solution with the lowest resulting elastic energy L is chosen to define the optimal Möbius transformation.

We have implemented the whole procedure into a Fortran program, MatchSurface.

4 EXPERIMENTAL RESULTS

4.1 Brain Surface Mapping

We demonstrate the feasibility of our algorithm by applying it to the brain surface matching problem. We consider a set of cortical surface models extracted from in-vivo MRI on 38 anonymous subjects [40]. These are the same models that were used by Yeo et al. to evaluate the performance of their




Fig. 3. *Automatic brain surface mapping.* The best registrations of the left and right hemispheres of the brain of Subject1 (Source) from the Anon database [40] onto the corresponding hemispheres of the template brain of Subject39 (Target) are shown on the inflated cortical surfaces of the target. The boundaries between the 35 regions obtained by manual parcellation (see text for details) are shown in red and green for the target and the images of the Source, respectively. Note that these are shown for illustrations only, as they were not used in the registration process.

Spherical Demons algorithm [41], with the exception that we removed Subject6, as its database of manual parcellation is incomplete. This data set covers significant anatomic variability as it contains young, middle aged, and elderly subjects and Alzheimer’s patients. For each subject, the left and right cortical surface models are provided as discrete genus-zero surface meshes, as well as information on their manual parcellation into 35 regions [42], as described in [41, Table II]. The left and right hemispheres of Subject39 are considered as templates (or Targets) on which all the other corresponding surfaces (Sources) are registered followed the procedure outlined in Algorithm 1; the optimization of the Möbius transformation from the unit sphere to itself was performed four times, each with a different initial guess derived from the principal components of the mesh (see implementation above). Usually, three of these optimizations led to local minima while the fourth one provided a good alignment between the template and source meshes, as illustrated in Fig. 3 for Subject1 in the database.

To assess the performance of our new algorithm, we repeated the Source to Target registrations described above with two widely used landmark-free registration methods as well as with a modified version of MatchSurface that uses landmarks. We first applied the FreeSurfer registration algorithm [23] using the default FreeSurfer settings. We also used the Spherical Demon algorithm [41]. As the latter only performs a registration of the sphere on itself, we used as input the spherical representations of the hemispheres generated by FreeSurfer. Finally, we used a modification of MatchSurface in which the landmark-free energy function $L(f)$ defined in (12) is replaced with a function $C(f)$ that computes the discrepancies between the positions p_i^T of landmark points on the Target and the images $f(p_i^S)$ of the corresponding points on the Source meshes:

$$C(f) = \sum_i (\|f(p_i^S) - p_i^T\|)^2. \quad (15)$$

Landmark points were set to sample uniformly the boundaries of all 35 regions. Each boundary was represented

with 70 to 500 points, for a total of 9,136 points. We note that all four methods include a mapping to the sphere.

To measure the quality of the cortical registrations provided by all four methods (MatchSurface, FreeSurfer, Spherical Demon, and the modified MatchSurface which we refer to as Landmark), we used a modified Hausdorff distances [43]. We projected each of the 35 regions from each of the Source subjects onto the Target surfaces and computed the mean distance between their boundaries. Fig. 4 displays the average distance per structure for the four algorithms, for the left and right hemispheres. Standard errors are shown as bars. The numbering of the structures correspond to [41, Table II]. Of the three landmark-free methods, the Spherical Demon algorithm performs best, followed by FreeSurfer, then by MatchSurface. The improved performance of Spherical Demon compared to FreeSurfer was already noticed [41]. Here we report that these two methods perform better than MatchSurface. We observe, however, that the differences between the three methods are small. It should be noted that FreeSurfer and Spherical Demons are designed to provide high-quality *local* alignments, while MatchSurface generates a *global* alignment of the two surfaces. It is therefore significant that a fully conformal correspondence gives landmark matching that is of comparable quality to nonconformal correspondences. The local versus nonlocal behavior is also significant to understand the results of the landmark-based method. Landmark performs significantly worse than the two local landmark-free methods, and slightly worse than MatchSurface. While Landmark is designed to obtain a good alignment of the landmarks, the optimization is performed globally in the space of conformal maps of the sphere to itself, i.e., in the space of Möbius transforms. Any such map can only align exactly three points between two representations of the sphere. When more landmarks are present, it provides a best-fit solution to the alignment problem with the constraint of imposing conformality, leading to some structures being poorly aligned. This limitation has already been described [34].

In Table 1, we compare the running times of the three landmark-free registration methods. MatchSurface is significantly faster for generating the mapping of a genus-zero surface onto the sphere as it solves this problem by minimizing a functional energy. MatchSurface is slower in generating the registration. We note, however, that our current implementation includes four minimizations of the elastic energy with four different initial conditions; we are working on eliminating this requirement, which would reduce the running time for registration by 4. The total running time of each method is obtained by summing the registration running time with twice the time needed for spherical mapping (as it is performed on the target and source meshes). MatchSurface is then found to be significantly faster than the two other methods.

4.2 Morphodynamics of Protein Structure Surfaces

Different experimental techniques provide different representations of protein structures. For example, high-resolution X-ray and NMR techniques generate atomistic models of proteins that are accurate at the Angstrom level, while techniques based on electron microscopy (EM) provide much




Fig. 4. Comparisons of MatchSurface, FreeSurfer, and Spherical Demons in registering cortical surfaces. The mean modified Hausdorff distances computed over the boundaries of the 35 structures are shown for (A) the left hemisphere, and (B) the right hemisphere.

lower resolution (typically 10-35 Angstroms) 3D electron density maps for the proteins of interest. While high resolution is preferred, the complexities of applying the corresponding techniques on large molecular systems and the comparative ease with which low-resolution EM and MS techniques can be used on the same system mean that low-resolution models are often available long before their high-resolution counterparts. Many techniques have been developed to achieve more details from EM and MS data by fitting atomic-resolution models into the low-resolution

density maps. These techniques work well if such models are available. However, such information is not available for many protein complexes, especially those that are dynamic. There is, therefore, a need to develop techniques that can analyze and measure the density maps directly. As these maps define surfaces, we can test the method described above and compare protein structures at a coarse level. We illustrate these tests on one protein, calmodulin. Results are shown in Fig. 5.

Calmodulin is a calcium binding protein expressed in all eukaryotic cells. It is a small protein that consists of two small domains separated by a linker region. The flexibility of this linker is key to the ability of calmodulin to bind to a wide range of substrates [44].

We consider two static structures for calmodulin corresponding to two different conformations, the apo (ligand-free) and holo (ligand-bound) forms, with PDB codes 1CLL and 1A29, respectively. We generated two trajectories between these two conformations that mimic the corresponding structural transition with MinActionPath. This program calculates the most probable trajectory between two known structural states, in the sense of maximum likelihood or minimum action [45]. The trajectory is described with 50 all-atom conformations. We then measured the evolution

TABLE 1
Running Times for Landmark-Free Registration

Method	Mapping on the sphere	Registration
FreeSurfer	104 (21) ^a	67.4 (11)
Spherical Demon	104 (21) ^b	2.15 (0.05)
MatchSurface	2.5 (0.8)	16.0 (3.0)

^a Mean running times and their standard deviations (in parenthesis) over all 74 registrations (37 subjects, 2 hemispheres). The running times, in minutes, were measured on a laptop computer using an Intel i7 CPU processor running at 1.8 GHz, with 4GB of RAM.

^b The Spherical Demon algorithm uses the spherical mapping generated by FreeSurfer.

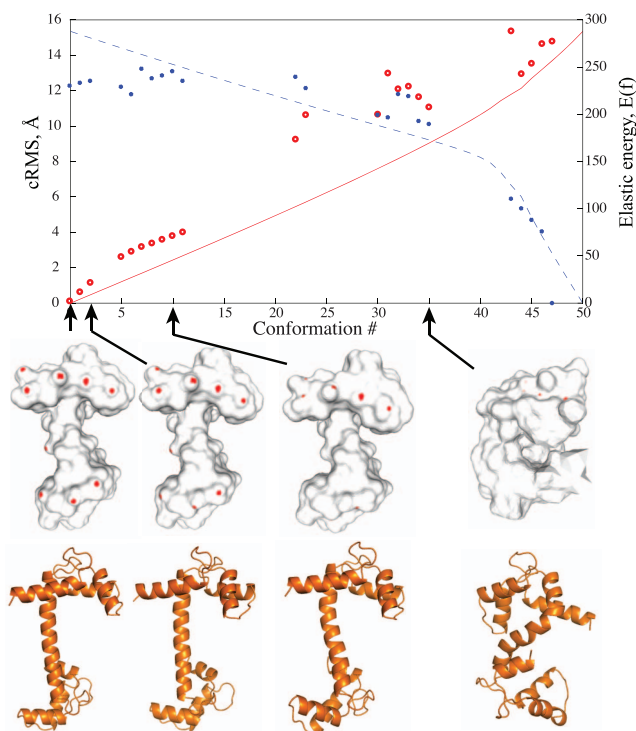


Fig. 5. Coarse and high-resolution analyses of the dynamics of calmodulin. We have generated a trajectory with 50 conformations between the ligand-free structure and a ligand-bound structure of calmodulin. This protein undergoes a considerable change of conformation as it transforms into the bound state, as is evident in the rightmost figure. The cRMS distances between the high-resolution structures of these conformers and the apo structure (conformation 0) are plotted versus the conformation number as a red solid line. We also plot the elastic energy between the surface meshes representing the same conformers and the surface mesh of the ligand-free protein as red dots; note that data are missing for some conformations as their skin meshes have genres greater than zero. We note a high level of correlation (0.96) between these two measures. The same behavior is observed when comparing the 50 conformations with the ligand-bound structure for cRMS (dashed blue line) and with the surface of conformer 47 for the surface-based comparison (blue dots); we could not use the surface of the ligand-bound structure as its genus is greater than 0. The surfaces reconstructed from the warped meshes and cartoon representations of the high-resolution structures are shown for a few conformations along the trajectory. The red dots on the protein surfaces illustrate the positions of landmark points manually picked on the ligand-free structure on the successive warped meshes.

of the structures of calmodulin along these trajectories by computing their coordinate root mean square deviations (cRMS) computed over C_α to the initial and to the final conformations. Results of these calculations are shown in Fig. 5 as solid and dashed lines, respectively. Note that these measures are based on knowledge at high resolution, as they are computed from the positions of the atoms.

In parallel, we compared the structures along the trajectories based on their skin surfaces [46]. We use the standard model in chemistry of representing a structure as a union of balls, with each ball corresponding to an atom. The skin surface of a protein is computed from the boundary of the union of these balls, where the center of a ball is given by the coordinates of the atom, and its radius is set to its van der Waals radius plus a probe radius of $Ra = 1.4 \text{ \AA}$. We generated quality meshes for the skin surfaces of all protein conformations along the trajectories for calmodulin using

the program smesh, described in detail in [47] and [48]. The corresponding triangular meshes have similar sizes for all proteins, with, on average, approximately 40,000 vertices and 70,000 triangles. We compared these meshes using MatchSurface. Results of these calculations are shown in Fig. 5. We note that it was not possible to perform the comparisons for all conformations, as for some their skin meshes have a genus larger than zero, and our method is currently limited to surfaces with genus zero. Higher genus is introduced when two nonadjacent patches of a protein surface come into contact and create a new handle. Geometric changes in the surface also occur when two adjacent patches fold together and no longer form part of the protein's boundary surface. While not changing the genus, this folding can cause abrupt changes in the elastic energy required to deform one surface onto another. We do observe that for all genus-zero conformations for which the comparisons were possible, the distances measured based on their surfaces (elastic energy, low resolution) correlate well with the distances measured based on the atomic coordinates (cRMS, high resolution), with correlation coefficients above 0.95.

4.3 MatchSurface Is Robust with Respect to the Quality of the Discrete Representations of the Surfaces

The implementations of the algorithms presented in Section 2 into the program MatchSurface were designed to be fast and robust. However, we do not control the qualities of the input meshes (both source and destination). To measure the sensitivity of MatchSurface to the mesh quality, we considered again the calmodulin protein and generated four different meshes for the same skin surface, from a fine mesh with 32,721 vertices and 65,438 triangles, to a much coarser mesh with 3,523 vertices and 7,042 triangles. The three coarse meshes were generated from the fine mesh with a procedure that maintains topology [49]. We then compared the skin surface with itself using MatchSurface, for all four possible source meshes and all four possible destination meshes, for a total of 16 comparisons. Results are shown in Fig. 6.

We note first that all elastic energies for the 16 comparisons are small, much smaller than the values recorded for different configurations of the surface (see Fig. 5). While this is expected as we are basically comparing a surface to itself, it remains a reassuring result as it illustrates the robustness of the method under change of mesh, as long as the mesh remains fine enough to accurately represent the geometry of the surface. The elastic energy can only be compared directly for the same source mesh, as it depends on the number of edges in the mesh, and increased elastic energy for finer meshes does not reflect increased geometric differences. For coarse source meshes, the quality of the destination mesh has little impact on the result. For a fine source mesh however, there is a loss of quality if the destination mesh is coarse. This is expected, as a coarse destination mesh provides a poor representation of its surface and the fine mesh will need to be distorted to fit on this coarse surface. This is seen in Fig. 6 for the case where the source mesh has 32,721 vertices, but not for the other cases.

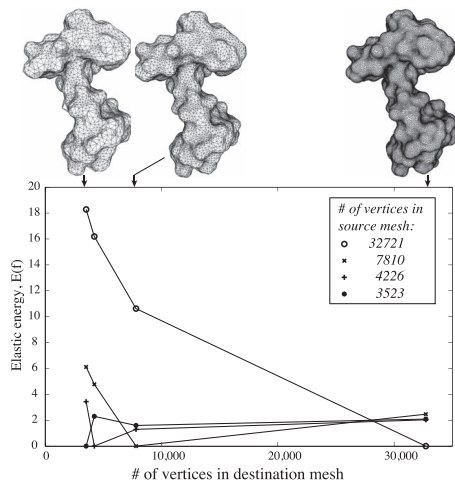


Fig. 6. *Sensitivity to mesh quality.* We performed all pairwise comparisons of four different discrete meshes representing the same skin surface of the protein calmodulin. Three of these meshes (two coarse and one fine) with 3,523, 7,810, and 32,721 vertices are shown on top. The corresponding elastic energies are plotted as a function of the number of vertices in the destination meshes, for the different source meshes. We note that all energy values are small (see Fig. 5 for comparison). This is expected as the underlying surface is the same; it is still encouraging as it illustrates robustness.

4.4 MatchSurface Is Sensitive to Small Variations in the Surfaces

To assess the sensitivity of MatchSurface to variations in the details of the surfaces, we generated different skin surfaces for the protein calmodulin by varying the probe radius Ra . A low probe radius defines a surface that follows closely the envelope of the union of atoms, also referred to as the vdW surface of the protein, while a large probe radius generates a much smoother surface, as all atoms have been significantly inflated (see Fig. 7). The skin surface obtained with $Ra = 1.5$ is the closest match to the accessible surface of the protein [50]; we use it as a reference. All the skin surfaces are then compared to this reference using MatchSurface. The resulting elastic energies are plotted against the probe radius Ra in Fig. 7. In these calculations, the reference mesh is used as a source, and warped around the different skin surfaces. Clearly, MatchSurface is very sensitive to the coarseness of the surface. The sensitivity is more pronounced as the source mesh is warped onto a surface with more details (see panel (B) of Fig. 7). We note that the differences measured between these different skin surfaces are much larger than the differences observed between different representations of the same surface (see Fig. 6).

5 SUMMARY AND CONCLUSIONS

We have developed a new method for automatically generating a conformal map between two surfaces of genus zero. This new approach leads to flexible registration of the two surfaces and accurate measurements of their geometric dissimilarities, without the need for the selection of landmark points. Its implementation within the program MatchSurface is based on fast and robust numerical methods, making surface comparisons feasible for a wide range of data sets. We have illustrated its use for brain surface mapping and protein surface comparisons. We have




Fig. 7. *Sensitivity to surface fluctuations.* (A) The elastic energy values obtained when comparing skin surfaces for calmodulin with the reference skin surface are plotted against the probe radius Ra used to generate these surfaces. We show three examples of these skin surfaces, for $Ra = 0.5$, 1.5 , and 3.0 Å. (B), (C), and (D) (bottom row) show the surfaces generated by the reference mesh after it has been warped on these three skin surfaces.

compared MatchSurface with FreeSurfer and Spherical Demon, two landmark-free methods for brain surface mappings. We have shown that a fully conformal correspondence generates surface registration that is of comparable quality to a non-conformal correspondence. We have demonstrated that the distances computed with our algorithm between low-resolution, surface-based representations of proteins are highly correlated with the corresponding distances computed between the high-resolution, atomistic models for the same proteins.

This paper, however, is just a first step toward achieving automatic, landmark-free registration of general surfaces. The current method has limitations that suggest direction for future work.

First, the method applies only to surfaces of genus zero. The discrete conformal mappings from the surfaces to the sphere rely on this property. In addition, we use the fact that a conformal self-mapping of the sphere belongs to the group of Möbius transformations, which provides significant simplification as a closed analytical form is available for members of that group. The concept of discrete conformal equivalence can be extended to surfaces with arbitrary topology, either through the introduction of cone singularities [31], or through the definition of a discrete conformal equivalence between a euclidean triangulation on the surface and a hyperbolic triangulation [36]. These alternate definitions lead to (discrete) mappings of the surface onto a domain in hyperbolic space (which can be represented by the Poincare disk model.) In general, there are no conformal maps between two surfaces of higher genus, though there are various constructions of




Fig. 8. *The problem of singularities.* We have used MatchSurface to compare two human models, one with one arm bent at ninety degrees (source), and the second with both arms extended and parallel to the body (target). The surface (S on T) reconstructed from the mesh obtained by warping the source mesh S on the target surface T is distorted, as it is missing the hands and feet. Vertices in these regions were assigned large, negative scaling factors u ; this is illustrated on the right panel that shows the surface S with all vertices whose factors u are lower than -10 colored in red.

closest-to-conformal mappings. Finding closest-to-isometric mappings for such surfaces is a topic for future studies.

Second, our algorithm works well for surfaces with uniform geometry with no extreme protrusions or spikes. This is usually the case for brain surfaces as well as for protein surfaces, but it is not true in general. In particular, MatchSurface does not perform well on models of humans, as illustrated on Fig. 8. The problem can be ascribed to the singularities introduced by the arms and legs of the models. Their discrete conformal mappings to the sphere introduce very large negative scaling factors on the vertices located at the hands and feet (see the left panel in Fig. 8, which in turn lead to infinitesimally small edge lengths in the projected metric \tilde{l} and consequently large numerical errors. This problem is not specific to our method, as it appears in many conformal mapping procedures. In some cases approximating by a conformal map appears to be too restrictive for accurate alignment. One solution is to introduce cone singularities in the regions with the worst distortions (see, for example, [31]). This brings us back to the first limitation discussed above and its possible resolution through the use of hyperbolic geometry.

Finally, our algorithm is limited to finding global matchings between a pair of surfaces of genus zero. A number of important applications would benefit from an extension that allows for partial matching. We are currently working on implementing such an extension.

ACKNOWLEDGMENTS

Patrice Koehl acknowledges support from the US National Institutes of Health. Joel Hass acknowledges support from the US National Science Foundation. The authors would like to thank Nina Amenta and Owen Carmichael for very useful discussions. They thank the authors of FreeSurfer and Spherical Demons for making their codes freely available.

REFERENCES

[1] R. Kötter and E. Wanke, "Mapping Brains without Coordinates," *Philosophical Trans. Royal Soc. of London, Series B, Biological Sciences*, vol. 360, pp. 751-766, 2005.

[2] A. Otte and U. Halsband, "Brain Imaging Tools in Neurosciences," *J. Physiology Paris*, vol. 99, pp. 281-292, 2006.

[3] A. Gholipour, N. Kehtarnavaz, R. Briggs, M. Devous, and K. Gonipath, "Brain Functional Localization: A Survey of Image Registration Techniques," *IEEE Trans. Medical Imaging*, vol. 26, no. 4, pp. 427-451, Apr. 2007.

[4] N. Max and E. Getzoff, "Spherical Harmonic Molecular Surfaces," *IEEE Computer Graphic and Applications*, vol. 8, no. 4, pp. 42-50, July 1988.

[5] P. Koehl, "Protein Structure Classification," *Reviews in Computational Chemistry*, K. Lipkowitz, T. Cundari, V. Gillet, and B. Boyd, eds., vol. 22, pp. 1-56, John Wiley & Sons, 2006.

[6] R. Kolodny, D. Petrey, and B. Honig, "Protein Structure Comparison: Implications for the Nature of Fold Space, and Structure and Function Prediction," *Current Opinion in Structural Biology*, vol. 16, pp. 393-398, 2006.

[7] V. Venkatraman, L. Sael, and D. Kihara, "Potential for Protein Surface Shape Analysis Using Spherical Harmonics and 3D Zernike Descriptors," *Cell Biochemistry and Biophysics*, vol. 54, pp. 23-32, 2009.

[8] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle, "Multiresolution Analysis of Arbitrary Meshes," *Proc. ACM SIGGRAPH '95*, pp. 175-182, 1995.

[9] P. Alliez, M. Meyer, and M. Desbrun, "Interactive Geometry Remeshing," *Proc. ACM SIGGRAPH '02*, pp. 347-354, 2002.

[10] Y. Lipman, R. Al-Aifari, and I. Daubechies, "The Continuous Procrustes Distance between Two Surfaces," *Communications in Pure and Applied Math.*, <http://arxiv.org/abs/arXiv:1106.4588v2> [math.DG], 2011.

[11] A. Bronstein, M. Bronstein, and R. Kimmel, "Generalized Multidimensional Scaling: A Framework for Isometry-Invariant Partial Surface Matching," *Proc. Nat'l Academy of Sciences USA*, vol. 103, pp. 1168-1172, 2006.

[12] F. Mémoli, "On the Use of Gromov-Hausdorff Distances for Shape Comparison," *Proc. Point Based Graphics*, pp. 81-90, 2007.

[13] Y. Lipman and I. Daubechies, "Conformal Wasserstein Distances: Comparing Surfaces in Polynomial Time," *Advances in Math.*, vol. 227, pp. 1047-1077, 2011.

[14] Y. Lipman, J. Puente, and I. Daubechies, "Conformal Wasserstein Distance: II. Computational Aspects and Extensions," *Math. of Computation*, <http://arxiv.org/abs/arXiv:1103.4681v2> [math.NA], 2011.

[15] D. Boyer, Y. Lipman, E. StClair, J. Puente, B. Patel, T. Funkhouser, J. Jernvall, and I. Daubechies, "Algorithms to Automatically Quantify the Geometric Similarity of Anatomical Surfaces," *Proc. Nat'l Academy of Sciences USA*, vol. 108, pp. 18221-18226, 2011.

[16] M. Zelditch, D. Swiderski, D. Sheets, and W. Fink, *Geometric Morphometrics for Biologists*. Elsevier Academic, 2004.

[17] H. Lu, L.-P. Nolte, and M. Reyes, "Interest Points Location for Brain Image Using Landmark-Annotated Atlas," *Int'l J. Imaging Systems Technology*, vol. 22, pp. 145-152, 2012.

[18] Q. Huang, B. Adams, M. Wicke, and L. Guibas, "Non-Rigid Registration under Isometric Deformations," *Proc. Symp. Geometry Processing*, pp. 1149-1458, 2008.

[19] R. Lasowski, A. Tevs, H.-P. Seidel, and M. Wand, "A Probabilistic Framework for Partial Intrinsic Symmetries in Geometric Data," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 963-970, 2009.

[20] R. Rustamov, "Laplace-Beltrami Eigenfunctions for Deformation Invariant Shape Representation," *Proc. Fifth Eurographics Symp. Geometry Processing*, pp. 225-233, 2007.

[21] J. Sun, M. Ovsjanikov, and L. Guibas, "A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion," *Proc. Symp. Geometry Processing*, pp. 1383-1392, 2009.

[22] B. Fischl, M. Sereno, and A. Dale, "Cortical Surface-Based Analysis. II: Inflation, Flattening, and a Surface-Based Coordinate System," *Neuroimage*, vol. 9, pp. 195-207, 1999.

[23] B. Fischl, M. Sereno, R. Tootell, and A. Dale, "High-Resolution Inter-Subject Averaging and a Coordinate System for the Cortical Surface," *Human Brain Mapping*, vol. 8, pp. 272-284, 1999.

[24] M. Valliant and J. Glaunès, "Surface Matching via Currents," *Lecture Notes in Computer Science*, vol. 3565, pp. 381-392, 2005.

[25] L. Bers, "Uniformization, Moduli, and Kleinian Groups," *Bull. London Math. Soc.*, vol. 4, pp. 257-300, 1972.

[26] J. Gray, "On the History of the Riemann Mapping Theorem," *Rendiconti del Circolo Matematico di Palermo*, ser. II, Supplemento 34, pp. 47-94, 1994.

- [27] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis, "Conformal Geometry and Brain Flattening," *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI '99)*, pp. 271-278, 1999.
- [28] M. Hurdal, P. Bowers, K. Stephenson, D. Sumners, K. Rehm, K. Shaper, and D. Rotenberg, "Quasiconformally Flat Mapping the Human Cerebellum," *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI '99)*, pp. 279-286, 1999.
- [29] X. Gu and S.-T. Yau, "Global Conformal Surface Parametrization," *Proc. Eurographics Symp. Geometry Processing*, pp. 127-137, 2003.
- [30] X. Gu, Y. Wang, T. Chan, P. Thompson, and S.-T. Yau, "Genus Zero Surface Conformal Mapping and Its Application to Brain Surface Mapping," *IEEE Trans. Medical Imaging*, vol. 23, no. 8, pp. 949-958, Aug. 2004.
- [31] B. Springborn, P. Schröder, and U. Pinkall, "Conformal Equivalence of Triangle Meshes," *Proc. SIGGRAPH Asia*, pp. 79-89, 2008.
- [32] Y. Lipman and T. Funkhouser, "Möbius Voting for Surface Correspondence," *ACM Trans. Graphics*, vol. 28, pp. 72-83, 2009.
- [33] D. Tosun, M. Rettmann, and J. Prince, "Mapping Techniques for Aligning Sulci across Multiple Brains," *Medical Image Analysis*, vol. 8, pp. 295-309, 2004.
- [34] Y. Wang, L. Lui, T. Chan, and P. Thompson, "Optimization of Brain Conformal Mapping with Landmarks," *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI '05)*, pp. 675-683, 2005.
- [35] A. Joshi, D. Shattuck, P. Thompson, and R. Leahy, "Surface-Constrained Volumetric Brain Registration Using Harmonic Mappings," *IEEE Trans. Medical Imaging*, vol. 26, no. 12, pp. 1657-1669, Dec. 2004.
- [36] A. Bobenko, U. Pinkall, and B. Springborn, "Discrete Conformal Maps and Ideal Hyperbolic Polyhedra," <http://arxiv.org/abs/arXiv:1005.2698> [math.GT], 2010.
- [37] B. Springborn, "A Unique Representation of Polyhedral Types. Centering via Möbius Transformations," *Math. Z.*, vol. 249, pp. 513-517, 2005.
- [38] C.-J. Lin and J. Moré, "Newton's Method for Large Bound-Constrained Optimization Problems," *SIAM J. Optimization*, vol. 9, pp. 1100-1127, 1999.
- [39] Y. Wu, Y. He, and H. Tian, "A Spherical Point Location Algorithm Based on Spherical Coordinates," *Proc. Int'l Conf. Computational Science and Its Applications*, vol. 3482, pp. 1099-1108, 2005.
- [40] A. Dale, B. Fischl, and M. Sereno, "Cortical Surface-Based Analysis. I: Segmentation and Surface Reconstruction," *Neuroimage*, vol. 9, pp. 179-194, 1999.
- [41] B. Yeo, M. Sabuncu, T. Vercauteren, N. Ayache, B. Fischl, and P. Golland, "Spherical Demons: Fast Diffeomorphic Landmark-Free Surface Registration," *IEEE Trans. Medical Imaging*, vol. 29, no. 3, pp. 650-668, Mar. 2010.
- [42] R. Desikan, F. Ségonne, B. Fischl, B. Quinn, B. Dickerson, D. Blacker, R. Buckner, A. Dale, R. Maguire, B. Hyman, M. Albert, and R. Killiany, "An Automated Labeling System for Subdividing the Human Cerebral Cortex on MRI Scans into Gyrus Based Regions of Interest," *Neuroimage*, vol. 31, pp. 968-980, 2006.
- [43] M.-P. Dubuisson and A. Jain, "A Modified Hausdorff Distance for Object Matching," *Proc. 12th IAPR Int'l Conf. Computer Vision and Image Processing*, pp. 566-568, 1994.
- [44] J. Chou, S. Li, C. Klee, and A. Bax, "Solution Structure of Ca(2)-Calmodulin Reveals Flexible Hand-Like Properties of Its Domains," *Nature Structural Biology*, vol. 8, pp. 990-997, 2001.
- [45] S.D.J. Franklin, P. Koehl, and M. Delarue, "MinActionPath: Maximum Likelihood Trajectory for Large-Scale Structural Transitions in a Coarse-Grained Locally Harmonic Energy Landscape," *Nucleic Acids Research*, vol. 35, pp. W477-W482, 2007.
- [46] H. Edelsbrunner, "Deformable Smooth Surface Design," *Discrete and Computational Geometry*, vol. 21, pp. 87-115, 1999.
- [47] H. Cheng and X. Shi, "Guaranteed Quality Triangulation of Molecular Skin Surfaces," *Proc. IEEE Visualization*, pp. 481-488, 2004.
- [48] H. Cheng and X. Shi, "Quality Mesh Generation for Molecular Skin Surfaces Using Restricted Union of Balls," *Proc. IEEE Visualization*, pp. 399-405, 2005.
- [49] X. Shi and P. Koehl, "Adaptive Surface Meshes Coarsening with Guaranteed Quality and Topology," *Proc. Int'l Conf. Computer Graphics*, pp. 53-61, 2009.
- [50] B. Lee and F.M. Richards, "Interpretation of Protein Structures: Estimation of Static Accessibility," *J. Molecular Biology*, vol. 55, pp. 379-400, 1971.



Patrice Koehl was trained as an engineer at the Ecole Centrale de Paris, France, where he graduated with a master's degree in bioengineering in 1984. He received the PhD degree in molecular biology from the University Louis Pasteur, Strasbourg, France, in 1989. He joined the French National Center for Research (CNRS) the same year as a senior scientist. In 1997, he visited the Department of Structural Biology at Stanford University; he extended his stay more than 7 years, becoming a senior research associate in that department. In 2004, he joined the University of California, Davis, where he is currently a professor of computer science and associate director of bioinformatics at its Genome Center. He is a recipient of the Bronze medal of the CNRS and is an Alfred P. Sloan fellow.



Joel Hass is currently a professor and chair of the Department of Mathematics, University of California, Davis, where he has been since 1988. Prior to that he held positions at the University of Michigan, Hebrew University of Jerusalem, and the University of California at Berkeley. He has held visiting positions at the University of Melbourne, Institute for Advanced Study in Princeton, and the Mathematical Sciences Research Institute in Berkeley. He is an Alfred P. Sloan fellow.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.