

# Globally Optimal Cortical Surface Matching With Exact Landmark Correspondence

Alex Tsui<sup>1</sup>, Devin Fenton<sup>1</sup>, Phong Vuong<sup>1</sup>, Joel Hass<sup>1</sup>, Patrice Koehl<sup>1</sup>, Nina Amenta<sup>1</sup>, David Coeurjolly<sup>2</sup>, Charles DeCarli<sup>1</sup>, and Owen Carmichael<sup>1</sup>

<sup>1</sup> University of California, Davis

<sup>2</sup> Université de Lyon, CNRS

**Abstract.** We present a method for establishing correspondences between human cortical surfaces that exactly matches the positions of given point landmarks, while attaining the global minimum of an objective function that quantifies how far the mapping deviates from conformality. On each surface, a conformal transformation is applied to the Euclidean distance metric, resulting in a hyperbolic metric with isolated *cone point* singularities at the landmarks. Equivalently, each surface is mapped to a hyperbolic *orbifold*: a pillow-like surface with each point landmark corresponding to a pillow corner. An initial surface-to-surface mapping exactly aligns the landmarks, and gradient descent is used to find the single, global minimum of the Dirichlet energy of the remainder of the mapping. Using a population of real MRI-based cortical surfaces with manually labeled sulcus endpoints as landmarks, we evaluate the approach by how much it distorts surfaces and by its biological plausibility: how well it aligns previously-unseen anatomical landmarks and by how well it promotes expected associations between cortical thickness and age. We show that, compared to a painstakingly-tuned approach that balances a tradeoff between minimizing landmark mismatch and Dirichlet energy, our method has similar biological plausibility, superior surface distortion, a better theoretical foundation, and fewer arbitrary parameters to tune. We also compare to conformal mapper in the spherical domain to show that sacrificing exact conformality of the mapping does not cause noticeable reductions in biological plausibility.

## 1 Introduction

Cortical surface matching—establishing point-wise correspondences between cerebral cortex surfaces—is a crucial step in MRI-based studies of brain morphology. Algorithms typically aim to induce a surface-to-surface mapping that minimally distorts morphological features. It is also desirable to use information provided by experts to guide the mapping. This information can consist of landmarks, labeled as points or curves on the surfaces, that are required to correspond to each other. Our goal is a cortical surface matching method that exactly matches point landmarks while insuring that the mapping minimizes distortion.

We quantify surface distortion in terms of *conformality*, or angle preservation. Conformal maps have been studied intensely due to their ability to preserve key

shape properties of biological specimens [9], and because conformal maps from any genus-zero surface to the sphere, and from any higher-genus surface to a surface of constant curvature, provably exist regardless of surface morphology [2] [21]. However, point landmarks are difficult to incorporate into conformal maps. A conformal map to the sphere, for example, is uniquely determined by the mapping of exactly three surface points; matching more than three requires sacrificing either conformality or exact landmark matching. For this reason, we seek maps that minimally deviate from conformality, using the Dirichlet energy of the mapping to quantify this deviation. Specifically, we conformally map each cortical surface to a hyperbolic orbifold, and find Dirichlet energy minimizing maps between the orbifolds that exactly align an arbitrary number of point landmarks. We show that the Dirichlet energy has exactly one, unique, global minimum over the relevant set of orbifold-to-orbifold maps, making it computationally robust.

Our approach is summarized in Figure 1. Given two triangulated cortical surfaces, we first alter the distance metrics on the two surfaces using conformal transformations [19]. This results in hyperbolic metrics on both surfaces with singularities at a finite number of isolated cone points, near which the metric behaves as though the surface is shaped locally like the vertex of a cone. One such cone point is located at each of the point landmarks. We then calculate an initial mapping from one surface to the other that exactly aligns the corresponding point landmarks. There is exactly one deformation of the initial map that maintains these point landmark matches while obtaining a minimum of the Dirichlet energy among the set of maps reachable by continuous deformation, *i.e.* within the homotopy class of the initial map. Given the unimodal nature of the Dirichlet energy landscape, finding the energy minimizing mapping within the homotopy class is straightforward using gradient descent.

Using cortical surfaces of elderly brains from a large-scale epidemiological study of aging [4] with manual point landmarks, we assessed whether our insistence on exact landmark matching (rather than approximate as in, *e.g.*, [15]) results in greater levels of surface distortion, and less usefulness in practical situations, that counterbalance the theoretical advantage of guaranteed globally-optimal mapping. We also assessed whether abandoning truly conformal mapping for Dirichlet energy minimizing mapping results in noteworthy practical disadvantages. To do this we compared our method (**OrbifoldExact**) to two competing methods that minimized landmark mismatch in a least-squares sense. An *orbifold least squares* method (**OrbifoldLS**), inspired by earlier work [15], balanced a tradeoff between Dirichlet energy minimization and landmark mismatch in a least squares sense. A *conformal least squares* method (**ConformalLS**) found the conformal map in the spherical domain that minimized landmark mismatch [15]. Experiments compared the methods in terms of point landmark mismatch, surface distortion, mismatch of novel (*i.e.*, not used to define the mapping) point landmarks, and ability to re-capitulate known population-level associations between cortical thickness and age [16]. Finally, experiments assessed whether the behavior of **OrbifoldLS** is stable with respect to critical but difficult-to-set operating parameters.

## 2 Prior work

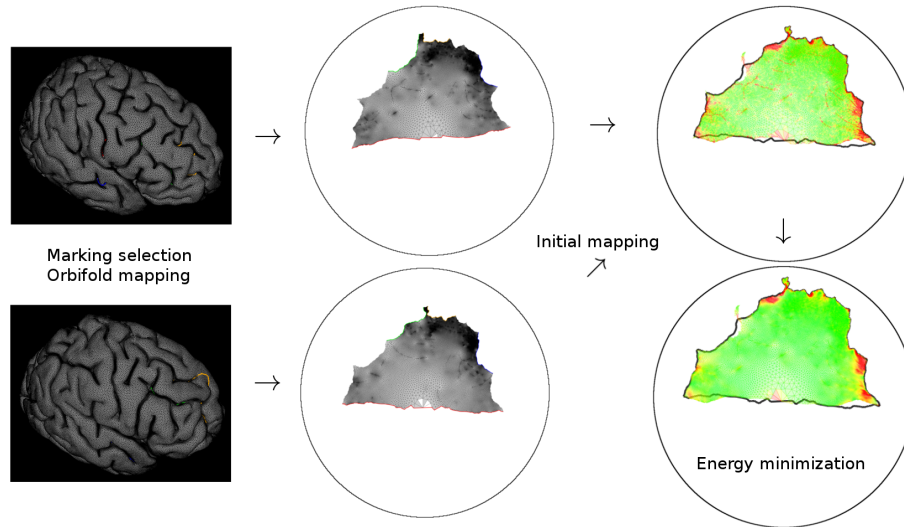
Prior work incorporating landmark information into cortical surface matching begins by finding harmonic energy minimizing mappings to spherical [15] or Euclidean [1] canonical domains as an initial step, or finding initial conformal maps to canonical Euclidean annuli [24] or hyperbolic "pairs of pants" [23]. One such canonical domain is then mapped onto the other in a way that encourages landmark matching. Using a Mobius transformation for this mapping [15] insures conformality but it is restricted to either exactly match only 3 points, or inexactly match a larger number. Harmonic maps are more flexible, but guarantee neither conformality nor exact landmark matching [1]. Quasi-conformal maps have bounded angle distortion, but the practical utility of very recently developed methods for computing angle-distortion-minimizing quasi-conformal maps is currently not clear [25]. Methods that cut the brain surface and map the cut to the boundary of a canonical domain have the additional limitation that conformality is lost along the cut [24], and there is an arbitrary decision about how exactly to map out the cut to the boundary. Note that while we specify paths between point landmarks that are similar to cuts, these instead constitute a *marking*, i.e. a landmark ordering convention that insures that the eventual mapping comes from a natural homotopy class, i.e. that the mapping can be connected to the desired optimal mapping by some deformation.

Building on earlier work on mapping cortical surfaces to the hyperbolic disc for visualization [9], we leverage an earlier observation that there is a single globally optimal map between hyperbolic discs that minimizes the Dirichlet energy [12], and model the landmarks as cone points in the hyperbolic metric to insure exact matching. Such cone point singularities have been considered previously to reduce area and length distortions during flattening to the plane, for applications such as texture mapping [22]. Other discrete conformal mapping methods [10] [19] incorporate cone point singularities into their conformal transformation of the surface distance metric without cutting.


## 3 Method

We begin with a pair of triangulated surfaces whose topology is spherical; in our experiments, each of these is the outer pial surface of a human cerebral cortex hemisphere output by commonly available software. Each surface has been annotated by an expert with a set of  $k$  point landmarks that are known to be in correspondence across surfaces. There are four key steps to our approach. The first, **orbifold mapping**, calculates conformally equivalent hyperbolic metrics on the surfaces, or equivalently, conformally maps the surface to a  $k$  pointed orbifold such that landmark points map to its cone points. Next, **marking** allows us to constrain the surface-to-surface mapping to a natural homotopy class of mappings: those that preclude reflections, complex surface folding, twisting, etc, in between exact point matches. For this step we define a tree that connects the landmark points analogously across surfaces. An **initial mapping** is constructed that maps the first tree to the second one, and extends that mapping

continuously across the rest of the surfaces; this mapping belongs to the natural homotopy class. Finally, **energy minimization** is used to adjust this mapping so that it arrives at the unique Dirichlet energy minimizing map within this homotopy class that leaves the landmark matches fixed. We outline each of these steps below.



**Fig. 1.** Overview of surface-to-surface mapping algorithm. Given two brain hemisphere surfaces (left), orbifold mapping constructs hyperbolic metrics on each surface. Then, marking selection identifies a natural homotopy class of mappings to optimize over. These steps allow us to map the surfaces to the hyperbolic plane (middle: sulci shown as colored boundary curves). Next, an initial surface-to-surface mapping is constructed in the hyperbolic plane (top right). This mapping introduces surface distortions in the form of dilatations (redder colors suggest greater distortion). Dirichlet energy minimization in hyperbolic space adjusts the mapping to obtain the global minimum of such distortions over the homotopy class (bottom right).



**Fig. 2.** Identifying a cortical surface (left) with a hyperbolic orbifold (right). Landmark points (white dots) are selected and connected by a set of paths to form a marking tree (colored curves, shown on an inflated cortical mesh, center). The point-to-point distance metric on the mesh can be conformally transformed into a hyperbolic metric by imposing an angle constraint at each landmark point, effectively identifying the brain with the hyperbolic orbifold.

### 3.1 Orbifold Mapping

Our method for calculating a conformal transformation of the surface distance metric is provided by Bobenko *et al.* [3]. Briefly, a surface triangulation  $T$  is a set of vertices  $V$ , edges  $E$ , and faces  $F$ , and a discrete metric  $l : V \rightarrow \mathbb{R}$  assigns lengths to each edge such that the triangle inequality at each face is satisfied. Two combinatorially equivalent triangulations, one with a Euclidean metric  $l$  and the other with hyperbolic metric  $\tilde{l}$ , are discretely conformally equivalent if their respective metrics  $l$  and  $\tilde{l}$  are related by

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

where  $u : V \rightarrow \mathbb{R}$  is an assignment of conformal factors to each vertex. This is an equivalence relation, and the set of discrete metrics in the same equivalence class is called a discrete conformal class. In this way, the conformal factors  $u$  define a conformal transformation of the Euclidean metric into a hyperbolic one.

Starting with a triangulation and the Euclidean metric  $(T, l)$ , we seek a discretely conformally equivalent triangulation with hyperbolic metric, such that the metric treats the  $k$  provided landmark points as cone points. One way to formulate this requirement is to consider the angle sum of vertex  $v_i \in V$ : this is the sum, over all mesh triangles that include  $v_i$ , of angles with  $v_i$  as the vertex. An angle sum of  $2\pi$  means that the surface is locally flat at  $v_i$ , while angle sums less than  $2\pi$  mean that the surface in the neighborhood of  $v_i$  more resembles the vertex of a cone. As described previously [3], we solve for conformal factors  $u$  such that the above equation is satisfied and each vertex is constrained to have a certain angle sum: for all point landmarks, the required sum is  $\pi$  to treat them as cone points; for all other points, the required sum is  $2\pi$  to treat them as locally flat. For a genus zero surface to have a hyperbolic orbifold metric with cone point singularities, the sum of the angle defects ( $2\pi$  - angle sum) across all vertices must be greater than  $4\pi$ . With an angle sum of  $\pi$  at the cone points and  $2\pi$  everywhere else, this means we must have at least five cone points to insure that the resulting metric is hyperbolic.

The energy function that is minimized to solve for the conformal factors  $u$ , as well as its first and second derivatives, are given in [3]. We use an implementation of a trust region version of Newton's method [13] to minimize this energy and compute the conformal factors.

### 3.2 Marking selection

Once the hyperbolic metric is defined on both surfaces, we impose an ordering on the set of landmarks that constrains the surface-to-surface mapping to be simple and well-behaved, *i.e.* to map corresponding points onto each other exactly while inducing no gross foldings or twists to the rest of the mapping. To do so, we first draw vertex-constrained paths that connect one sulcal endpoint to the opposite endpoint. We then insert additional paths connecting endpoints across sulci curves until all sulcal endpoints are connected in a tree structure,

called the *marking tree* (see Figure 2 middle). Note that individual paths must not self-intersect and pairs of paths must not intersect each other, but otherwise any tree of such paths will suffice.

The significance of the marking tree is that it specifies the homotopy class of the mapping. That is, it allows us to specify an initial mapping such that one tree is mapped onto the other in a natural way, and the tree-to-tree mapping is extrapolated out continuously over the rest of the surfaces. Optimization is then constrained such that only mappings within the same homotopy class of this initial mapping are allowed. Thus, this step effectively allows us to rule out unnatural mappings. Together with the fact that under hyperbolic space, there exists a unique harmonic map that minimizes Dirichlet energy in each homotopy class, the mapping can be refined to obtain the globally unique map in the sense of minimum Dirichlet energy.

### 3.3 Initial Mapping

Because there is a single, unique global minimum of the Dirichlet energy within each homotopy class, the initial mapping that is subsequently optimized to minimize the Dirichlet energy is highly arbitrary— the only requirement is that it belong to the homotopy class of mappings that map one surface onto the other in a simple, reasonable way (*i.e.*, with no complex folding or twisting of the surface in between landmarks). We parameterize this initial mapping in the Euclidean topological disc to conveniently make use of existing computational methods [12] that insure the mapping is in the correct homotopy class, although we emphasize that the optimization of the mapping itself is governed by the hyperbolic metric induced upon it as described above. The initial mapping is constructed by assigning each edge of the marking tree to a side of a regular polygon in the Euclidean disc, and filling in the remainder of the mapping by minimizing its harmonic energy, with the marking tree edges as a fixed boundary. One such harmonic energy minimization is performed per surfaces; the overlay of the two Euclidean polygons provides the initial surface-to-surface mapping. The corners of the polygon are identified with the landmark points in the order that they appear in a traversal around the marking tree, and thus the initial mapping maps one marking tree to the other in a simple way while filling in the remainder of the mapping in a smooth, reasonable manner. We emphasize that while this approach effectively cuts the spherical-topology surface open along the marking tree, resulting in a topological disc in 3D that is then flattened into a disc contained in the plane as in prior work [12], this is solely for the purpose of establishing an approximate initial mapping that is then optimized based on the hyperbolic metrics described above; the optimized mapping does not contain discontinuities or other distorting artifacts along the marking tree edges. We also emphasize that this initial mapping is highly arbitrary: a variety of simpler alternative methods may be applicable and the specific method does not impact the optimality of the final surface-to-surface mapping.

### 3.4 Energy minimization

Suppose one surface contains edges  $e_{ij}$  that connect vertices  $v_i$  to  $v_j$ ,  $w_{ij}$  are the cotangent weights  $w_{ij} = 0.5(\cot \alpha + \cot \beta)$ , where  $\alpha$  and  $\beta$  are the two angles opposite the edge  $e_{ij}$ , and the other surface lies in the hyperbolic plane. The initial mapping  $f$  maps the points of this surface onto the other surface such that point landmarks are kept in correspondence. The Dirichlet energy of  $f$  under the hyperbolic metric can be approximated as follows:

$$E(f) = \frac{1}{2} \sum_{e_{ij}} w_{ij} (f(v_j) - f(v_i))^2 \quad (2)$$

Given hyperbolic orbifold structures defined on each surface and an initial map between them, the theorems of Eells and Sampson [5] and of Hartman [8] imply that there is a unique harmonic map that minimizes Dirichlet energy in the homotopy class of maps that can be realized by deforming the original map.

To compute the Dirichlet energy minimizing map, we re-parameterize the the initial 2D polygon-to-polygon mapping in the Euclidean disc to the hyperbolic disc (specifically the Poincare disc). We use steepest descent to minimize the Dirichlet energy: this amounts to adjusting surface vertex positions in the Poincare disc but constraining the vertices corresponding to landmark points to stay fixed and in correspondence. To overcome numerical issues, we follow a prior approach by [12] optimizing the mapping of each surface vertex one at a time: this point and its surrounding one-ring of mesh faces (its local "chart") is translated to the Poincare disc origin, where the hyperbolic metric is well approximated by a corresponding Euclidean metric. The position of that point is then adjusted to minimize the Dirichlet energy and translated back to its original position. See [20] for additional details on the construction and management of the local charts.

## 4 Experiments

### 4.1 Data

We obtained brain MRI of 50 healthy elderly subjects from a prior study [4], identified gray matter voxels [6], used BrainVisa to convert each hemisphere's cortical gray matter mask into matching inner and outer pial surface meshes [14], from which we removed small or slivery mesh triangles [7]. Cortical thickness was estimated at each outer pial surface vertex using a "normal-average" approach [11]. A set of 16 sulcal endpoints were annotated on each outer pial hemisphere by an expert rater using a validated protocol [18].

### 4.2 Competing methods

Experiments compared the method described above, termed **OrbifoldExact**, to two competing methods that strike a different balance between landmark

matching and surface distortion. One competitor, `OrbifoldLS`, is identical to `OrbifoldExact` except that point landmarks are not constrained to be fixed during energy minimization, and the energy function balances a tradeoff between landmark mismatch and Dirichlet energy:

$$E(f) = (1 - \lambda) \frac{1}{2} \sum_{e_{ij}} w_{ij} (f(v_j) - f(v_i))^2 + \lambda \sum_{v_i \in L_1} (f(v_i) - L_2(v_i))^2 \quad (3)$$

where  $L_1$  is the set of landmark points on the source surface and  $L_2(v_i)$  is the landmark point on the target surface corresponding to landmark point  $v_i$  on the source surface.

The other competitor, `ConformalLS`, first conformally maps each triangular mesh onto the unit sphere [19], then solves for a Möbius transformation (*i.e.*, a conformal mapping of the first sphere to the second one) that minimizes point landmark mismatch in the same least squares sense as in `OrbifoldLS`.

### 4.3 Performance measures

**Surface distortion.** `OrbifoldLS` and `OrbifoldExact`, are able to induce surface distortions in the form of *dilatations*—stretches that transform local circles to local ellipses under the mapping—while `ConformalLS` precludes such dilatations by construction. For the former, we compute a discrete approximation of dilatation [17] at every mesh triangle and report summaries of dilatation over all vertices. We also show brain surfaces color-coded by dilatation under various mappings.

**Landmark mismatch.** `OrbifoldLS` and `ConformalLS` allow imperfect matching of point landmarks, while `OrbifoldExact` requires exact landmark matches by construction. We report the mean Euclidean distance between corresponding landmark points under the mapping for the former two approaches.

**Strength of expected associations.** We selected one of the 50 left outer pial surfaces as a canonical brain surface and used each of the three techniques to map the remaining 49 surfaces onto it. The mappings allowed us to transfer cortical thicknesses from the 49 surfaces onto the canonical one, and interpolate the thicknesses to the positions of canonical mesh points. This resulted in 50 cortical thicknesses (one per subject) at each canonical mesh point. We calculated a linear regression model at each mesh point to assess the strength of association between that point’s local cortical thickness and the age of the corresponding subjects. The  $p$  values for these regressions were corrected for multiple comparisons [26], and the  $p$  values at the vertices were interpolated across intervening mesh faces. The surface area that had  $p < .05$  was then calculated. Numerous studies (*e.g.*, [16]), agree that the thickness of the cortical mantle reduce with age, so we seek mapping methods that give rise to a statistically significant relationship with age across the largest possible cortical area.



#### 4.4 Experimental settings

**Comparison of 3 methods.** Given 16 point landmarks on a hemisphere surface, we consider two experimental settings. In the first, we use the full set of landmarks to define the mappings, and evaluate surface distortion, landmark mismatch, and strength of expected associations on appropriate methods. In the second, we cross-validate: we use 14 of the 16 landmarks to define the mapping and evaluate landmark mismatch for the remaining two.

**Practical limitations of OrbifoldLS.** OrbifoldLS includes two operating parameters that are difficult for a user to know how to set optimally: the marking tree and  $\lambda$ . We assessed whether settings of these parameters impact landmark mismatch and surface distortion by running OrbifoldLS over a range of settings for both and assessing variability in both performance characteristics.


## 5 Results

Landmark mismatch statistics for OrbifoldLS and ConformalLS are shown in Table 1. Mismatch is substantial for each method, averaging greater than 5 *mm* for most sulci (consider that the mean mesh edge length is 0.55 *mm*). This motivates exact landmark matching as in our method: the mapping problem is difficult enough that least squares methods do not simply happen find a solution that easily matches landmarks close to exactly. Dilatation statistics for OrbifoldLS and OrbifoldExact are listed in Table 2, and anecdotal dilatation maps are shown in Figure 3. As expected, OrbifoldExact gives rise to dramatic maximum dilatations, which occur at the isolated point landmarks where the orbifold construction has dramatically changed the surface distance metric. But large dilatations do not broadly infect large surface regions (see Figure 3 and mean dilatations in Table 2), suggesting that requiring exact landmark matches does not preclude broadly well-behaved mappings.

Sulcus	ConformalLS	OrbifoldLS, $\lambda = 0.1$	OrbifoldLS, $\lambda = 0.5$
Central	19.95	7.62	6.29
Precentral	20.01	7.80	6.33
Postcentral	21.42	9.56	7.40
Cingulate	26.61	7.16	6.11
Intraparietal	25.54	8.66	7.25
Superior Temporal	29.40	8.16	6.46
Superior Frontal	24.12	6.45	5.28
Inferior Frontal	24.31	5.90	4.90

**Table 1.** Average landmark mismatches, in *mm*, for sulcal endpoints across 50 subjects. Note that landmark mismatch for OrbifoldExact is zero by construction.

Landmark mismatch statistics for the left-out landmark in the cross-validation setting are shown in Table 3. OrbifoldLS conferred no notable benefit over OrbifoldExact, again suggesting that requiring exact landmark matches draws no appreciable cost in terms of practical performance. ConformalLS provides superior matching of the left-out point for approximately half of the landmarks, but matching is similar to OrbifoldExact or worse for the other half. This



**Fig. 3.** Mesh face dilations for one example pair of surfaces (greener/redder indicates lesser/greater dilatation). Recall that dilatation is exactly one for ConformalLS. Left to right, columns show dilatation of the initial mapping, OrbifoldExact, OrbifoldLS with  $\lambda = 0.5$ , and OrbifoldLS with  $\lambda = 0.1$ . White dots indicate landmark points.

Method	Mean	Median	Max	Method	Age
OrbifoldExact	1.489	1.300	2089	ConformalLS	450.0
OrbifoldLS, $\lambda = 0.1$	1.452	1.283	320	OrbifoldExact	3885.8
OrbifoldLS, $\lambda = 0.5$	1.503	1.288	732	OrbifoldLS, $\lambda = 0.1$	4020.8

**Table 2. Left:** Mean, median, and maximum distortion from conformality (dilatation) across all mesh faces on all 50 surfaces, for three competing methods. Note that dilatation of 1 is realized in an isometry and by construction is 1 everywhere for ConformalLS. **Right:** Total area ( $mm^2$ ) of cortex showing statistically significant evidence of a linear relationship with cortical thickness.

suggests that ConformalLS offers no obvious performance advantage over OrbifoldLS in terms of matching previously-unseen landmarks, and the landmarks used to define the mapping again exhibit substantial mismatch.

Table 2 summarizes the cortical surface area significantly associated with age. OrbifoldExact and OrbifoldLS are highly similar in this regard, while ConformalLS lags far behind. This suggests that OrbifoldLS and ConformalLS holds no obvious practical advantage in terms of relevant applications that might compensate for their other theoretical or computational limitations.

Sulcus	ConformalLS	OrbifoldLS, $\lambda = 0.1$	OrbifoldLS, $\lambda = 0.5$	OrbifoldExact
Central	12.49 (20.90)	25.59 (7.22)	25.61 (6.53)	25.18
Cingulate	28.96 (21.23)	26.65 (6.93)	26.65 (6.27)	25.20
Inferior Frontal	24.31 (21.11)	24.57 (6.87)	24.56 (6.03)	22.82
Intraparietal	30.44 (21.21)	31.37 (7.08)	31.35 (6.18)	30.39
Postcentral	21.62 (20.28)	28.92 (7.01)	28.92 (6.20)	30.93
Precentral	17.04 (21.05)	25.84 (6.83)	25.85 (6.01)	26.12
Superior Frontal	23.86 (20.77)	36.75 (5.79)	35.89 (4.95)	35.90
Superior Temporal	39.57 (23.52)	33.84 (7.69)	33.85 (6.81)	31.16

**Table 3.** Mean sulcal endpoint deviations in  $mm$  of the left-out sulcus using a map created by applying the given method. Numbers in parentheses indicate the landmark deviation ( $mm$ ) averaged over the remaining seven sulci.

Variability in landmark mismatch for `OrbifoldLS` with respect to marking tree and  $\lambda$  are summarized in Table 4. Landmark matching varies substantially with respect to these two parameters, although users have no way of gleaned that one or another setting is preferable *a priori*. These results suggest that, while the matching performance of `OrbifoldLS` and `OrbifoldExact` are similar, variability due to arbitrary parameter settings reduces the robustness of `OrbifoldLS`. Indeed, this demonstrates the difficulty in least squares approaches as optimizing an energy function with local minima is subject to initial conditions.

Marking tree	Deviation ( <i>mm</i> )	$\lambda$	Deviation ( <i>mm</i> )
A	2.05	0.1	2.05
B	0.23	0.5	0.24
C	0.53	0.8	0.027

**Table 4.** Mean landmark deviations with respect to marking tree and  $\lambda$ . Experiment was conducted with eight point landmarks. We set  $\lambda = 0.1$  when varying the marking tree, and we fix marking tree A when varying  $\lambda$ . Tree A is the same as shown in Figure 2. Trees B and C connect the landmarks instead in a single path.

In conclusion, our method provides a surface-to-surface mapping that exactly matches point landmarks and arrives at the global minimum of a particular surface distortion energy. Experiments suggest that neither requiring exact landmark matches, nor failing to require conformality, reduce the practical performance of the method, suggesting usefulness in practice.

## References

1. G. Auzias, J. Lefèvre, A. Le Troter, C. Fischer, M. Perrot, J. Régis, and O. Coulon. Model-driven harmonic parameterization of the cortical surface. In *Proceedings of the 14th international conference on Medical image computing and computer-assisted intervention, MICCAI'11*, pages 310–317, 2011.
2. L. Bers. Uniformization, moduli, and kleinian groups. *Bull. London Math. Soc.*, 4:257–300, 1972.
3. A. Bobenko, U. Pinkall, and B. Springborn. Discrete conformal maps and ideal hyperbolic polyhedra. pages 1–49, 2010.
4. S. Debette, A. Beiser, C. DeCarli, R. Au, J.J. Himali, M. Kelly-Hayes, J.R. Romero, C.S. Kase, P.A. Wolf, and S. Seshadri. Association of mri markers of vascular brain injury with incident stroke, mild cognitive impairment, dementia, and mortality the framingham offspring study. *Stroke*, 41(4):600–606, 2010.
5. J. Eells and J.H. Sampson. Harmonic mappings of riemannian manifolds. *American journal of mathematics*, 86:109–160, 1964.
6. E. Fletcher, B. Singh, D. Harvey, O. Carmichael, and C. DeCarli. Adaptive image segmentation for robust measurement of longitudinal brain tissue change. In *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pages 5319–5322. IEEE, 2012.
7. S. Fuhrmann, J. Ackermann, T. Kalbe, and M. Goesele. Direct Resampling for Isotropic Surface Remeshing. In *Vision, Modeling, and Visualization –VMV '10*, pages 9–16, 2010.

8. P. Hartman. On homotopic harmonic maps. *Canadian journal of mathematics*, 19:673–687, 1967.
9. M.K. Hurdal and K. Stephenson. Discrete conformal methods for cortical brain flattening. *NeuroImage*, 45:S86–S98, 2009.
10. L. Kharevych, B. Springborn, and P. Schröder. Discrete conformal mappings via circle patterns. *ACM Trans. Graph.*, 25:412–438, 2006.
11. P. Kochunov, W. Rogers, J.-F. Mangin, and J. Lancaster. A library of cortical morphology analysis tools to study development, aging and genetics of cerebral cortex. *Neuroinformatics*, 10:81–96, 2012.
12. X. Li, Y. Bao, X. Guo, M. Jin, X. Gu, and H. Qin. Globally optimal surface mapping for surfaces with arbitrary topology. *IEEE Trans. Vis. Comput. Graphics*, 14:805–19, 2008.
13. C. Lin and J. Moré. Newton’s method for large bound-constrained optimization problems. *SIAM Journal on Optimization*, 9:1100–1127, 1999.
14. J.F. Mangin, V. Frouin, I. Bloch, J. Régis, and J. López-Krahe. From 3d magnetic resonance images to structural representations of the cortex topography using topology preserving deformations. *Journal of Mathematical Imaging and Vision*, 5(4):297–318, 1995.
15. L. Ming, Y. Wang, T.F. Chan, and P. Thompson. Landmark constrained genus zero surface conformal mapping and its application to brain mapping research. *Appl. Numer. Math.*, 57:847–858, 2007.
16. D.H. Salat, R.L. Buckner, A.Z. Snyder, D.N. Greve, R.S.R. Desikan, E. Busa, J.C. Morris, A.M. Dale, and B. Fischl. Thinning of the cerebral cortex in aging. *Cerebral Cortex*, 14(7):721–730, 2004.
17. P.V. Sander, J. Snyder, S.J. Gortler, and H. Hoppe. Texture mapping progressive meshes. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques - SIGGRAPH '01*, pages 409–416, 2001.
18. E.R. Sowell, P.M. Thompson, D. Rex, D. Kornsand, K.D. Tessner, T.L. Jernigan, and A.W. Toga. Mapping sulcal pattern asymmetry and local cortical surface gray matter distribution in vivo: maturation in perisylvian cortices. *Cerebral cortex*, 12:17–26, 2002.
19. B. Springborn, P. Schröder, and U. Pinkall. Conformal equivalence of triangle meshes. *ACM Trans. Graph.*, 27:1–11, 2008.
20. V. Surazhsky and C. Gotsman. Explicit surface remeshing. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, SGP '03, pages 20–30, 2003.
21. W.P. Thurston. Three dimensional manifolds, kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc.*, 6:357–382, 1982.
22. Y. Tong, P. Alliez, D. Cohen-Steiner, and M. Desbrun. Designing quadrangulations with discrete harmonic forms. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, SGP '06, pages 201–210, 2006.
23. Y. Wang, W. Dai, Y.-Y. Chou, X. Gu, T.F. Chan, A.W. Toga, and P.M. Thompson. Studying brain morphometry using conformal equivalence class. In *IEEE 12th International Conference on Computer Vision, ICCV '09*, pages 2365–2372, 2009.
24. Y. Wang, X. Gu, T.F. Chan, P.M. Thompson, and S.-T. Yau. Conformal slit mapping and its applications to brain surface parameterization. In *Proceedings of the 11th international conference on Medical Image Computing and Computer-Assisted Intervention, MICCAI '08*, pages 585–593, 2008.
25. O. Weber, A. Myles, and D. Zorin. Computing extremal quasiconformal maps. *Computer Graphics Forum*, 31:1679–1689, 2012.

26. K.J. Worsley, S. Marrett, P. Neelin, A.C. Vandal, K.J. Friston, A.C. Evans, et al. A unified statistical approach for determining significant signals in images of cerebral activation. *Human brain mapping*, 4(1):58–73, 1996.