Area-preserving Surface Flattening
Using Lie Advection

Guangyu Zou[1], Jiaxi Hu[1], Xianfeng Gu[2], and Jing Hua[1]

[1] Wayne State University, USA
[2] State University of New York at Stony Brook, USA

Abstract. In this paper, we propose a novel area-preserving surface flattening method, which is rigorous in theory, efficient in computation, yet general in application domains. Leveraged on the state-of-the-art flattening techniques, an infinitesimal area restoring diffeomorphic flow is constructed as a Lie advection of differential 2-forms on the manifold, which yields strict equality of area elements between the flattened and the original surfaces at its final state. With a surface represented by a triangular mesh, we present how an deterministic algorithm can be faithfully implemented to its continuous counterpart. To demonstrate the utility of this method, we have applied our method to both the cortical hemisphere and the entire cortex. Highly complied results are obtained in a matter of seconds.

Keywords: Brain mapping, area-preserving flattening, Lie advection

1   Introduction

Given the fact that many anatomical surfaces are intrinsically 2D and highly undulated, flattening techniques constitute a major means of visualizing pathologies that are deeply buried within the folds [4]. Flattening (parametrization) also enables many procedures on a regular parameter domain, therefore resulting in more efficient and stable computation. However, existing flattening methods, including conformal mapping [2, 6, 5], typically suffer from severe, unpredictable area distortion when dealing with extruding shapes that contains rich, complex features, which largely impedes the capture of anatomical characteristics as well as other associated imaging modalities on a planar domain. Although it is known that surface cuts can effectively reduce the distortion of final flattening, such practice is typically not preferred, since neighborhoods on opposite sides of a cut will become far apart in the final flattened representation.

The technique presented in this paper theoretically ensures a uniform sampling of the surface on the parameter domain. Each patch retains exactly the same area when flattened to 2D. Based on the mathematical advance, a series of analysis tasks regarding neuronal density, activation extent, cortical thickness, among others, can be instead invoked in 2D where more compact data structures, more efficient discretization schemes and faster data access are available. Towards the same goal, a handful of respectable efforts have been made [3, 4, 8,

9]. However, the flattening results typically correspond only to local minima with respect to certain objective functionals [3, 4], and lack provable guarantee of area preservation. Pons et al. [8] designed a tangential motion according to a given normal motion in a cortical inflating procedure, which preserves area exactly, but is not capable of flattening a 3D surface into 2D. A mass (e.g., area/volume) preserving mapping was explicitly sought in $\mathbf{R}^n$ by a gradient decent method to the Monge-Kantorovich functional in [9], implemented on a regular Cartesian grid. It is not clear how this method can be extended to a general manifold.

In contrast, our method computes global strict area-preserving (A.P.) flattening of arbitrary 2-manifolds using *Lie advection*, a concept from classical mechanics. To our best knowledge, this is the first work that employs Lie advection as a tool to manipulate area changes in the context of surface flattening. Besides a general framework, our method also allows an efficient, yet accurate discretization scheme that is motivated by preserving the original geometric and algebraic structures of the continuous model in the limit, therefore rendering better numerical fidelity.

A similar idea was mathematically sketched in [1]. The discussion was restricted to a spherical domain. Despite similar concepts, our method is deterministic and derived for arbitrary surface manifolds. In the remainder of this paper, we take brain surface as an acting example, while the presented method is principally applicable to general surfaces.

2   Basic Idea

Starting with an arbitrary initial diffeomorphism from an given surface to the desired domain, e.g., a conformal parametrization [7, 5], we can subsequently evolve it to an A.P. alternative as follows. Suppose M and N are two differentiable 2-manifolds, associated by a diffeomorphism $f : M \rightarrow N$. The area element of a surface is a differential 2-form. Let ω_i, $i = M, N$ be the area form of M and N, respectively. The pullback of ω_N under f is a differential 2-form on M, denoted as $f^*(\omega_N)$. Suppose M and N have the same area integral after an appropriate scaling, that is, $\int_M \omega_M = \int_N \omega_N$. Computing an A.P. map $\mu : M \rightarrow N$ now is equivalent to finding a diffeomorphism $\varphi : M \rightarrow M$, such that $\varphi^*(\omega_M) = f^*(\omega_N)$. Therefore, μ is given by $f \circ \varphi^{-1}$. To accomplish this, we first linearly interpolate a 2-form over time:

$$\omega_t = (1 - t)\omega_M + tf^*(\omega_N), \quad t \in [0, 1]. \tag{1}$$

Note that $\omega_0 = \omega_M$ and $\omega_1 = f^*(\omega_N)$. In the following, we will design a one parameter family of diffeomorphisms, such that the corresponding flow deforms the area element in the same fashion as ω_t.

More specifically, consider a smooth surface M with a smooth vector field V on it. Given any point $p \in M$, there exist a unique integral curve $\gamma(t)$ of V passing through it, such that

$$\begin{cases} \frac{d\gamma_p(t)}{dt} = V(\gamma_p(t)), \\ \gamma_p(0) = p. \end{cases} \tag{2}$$

A one parameter family of diffeomorphisms (which are also automorphisms) ϕ_t, parameterized by $t \in [0, 1]$, can be defined on M as

$$\phi_t(p) = \gamma_p(t). \tag{3}$$

We want $\phi_t^*(\omega_0) = \omega_t$. Substituting p with ω_t in Eq. (3) and computing time derivative at $t = 0$, we get

$$\frac{d\phi_t^*\omega}{dt}\bigg|_{t=0} = \omega_M - f^*(\omega_N), \tag{4}$$

which, by definition, is the *Lie derivative* of ω_t with respect to V. Hence, the central equation to solve is

$$\mathcal{L}_{V(t)}\omega_t = \omega_M - f^*(\omega_N), \tag{5}$$

where $\mathcal{L}_{V(t)}$ denotes the Lie derivative with respect to $V(t)$. Intuitively, $\mathcal{L}_{V(t)}$ estimates the change of ω_t along the flow of $V(t)$. By using Cartan's formula $\mathcal{L}_V = di_V + i_V d$, where i_V denotes the interior product with respect to $V(t)$, we have

$$d(i_V\omega_t) + i_V d\omega_t = \omega_M - f^*(\omega_N). \tag{6}$$

Since ω_t is a 2-form, $d\omega_t = 0$ on M. Hence, we have

$$d(i_V\omega_t) = \omega_M - f^*(\omega_N). \tag{7}$$

By definition, we can write $\omega_M = h_1 du \wedge dv$ and $f^*(\omega_N) = h_2 du \wedge dv$. h_1 and h_2 are the scaling factors from a mutual parameter domain to M and N, respectively. Now let

$$\Delta g = h_2 - h_1, \tag{8}$$

where Δ denotes the Laplacian-Beltrami differential operator. We substitute Eq. 8 into Eq. 7. Using the fact that M and N are 2-manifolds, $V(t)$ can be solved as

$$V(t) = \frac{1}{(1-t)h_1 + th_2}\nabla g. \tag{9}$$

Note that, as implied by Eq. (9), $V(t)$ varies both in magnitude and direction over time. we need to solve $V(t)$ at each time. g is essentially a harmonic scalar field on M. As a result, the corresponding gradient vector field ∇g is guaranteed by construction to be highly smooth and free of extraneous critical points. Time integration of $V(t)$ yields a diffeomorphism. Finally, the A.P. mapping is given by $f \circ \phi_{t=1}^{-1} : M \to N$.

3   Algorithm

In practice, the surface is represented by a triangular mesh. The position of vertex v_i is denoted by $\mathbf{v}_i \in \mathbf{R}^3$. The 1-ring neighbors of v_i is denoted as $N_1(i)$.

For a vertex v_i, the associated surface patch is chosen to be the barycentric finite volume, denoted as A_i.

We first assume the target domain is a unit square, which is commonly used to parameterize topological disks. Later, we will show how this framework can be extended to the unit sphere to flatten closed genus zero surfaces in spherical geometry. The square boundary condition is set up as follows: first, the boundary is isometrically mapped to that of the unit square D; next, a discrete conformal map of the interior is computed in the least squares sense [7]. Fig. 1(a) and 1(b) show the lateral and mesial views of a cortical hemisphere, respectively. The initial (conformal) flattening is illustrated in Fig. 1(c), the area distortion of which is color encoded in Fig. 1(d). Notice how some lateral cortical patterns suffer from the intense geometric stretch, which greatly impairs inspection of these regions. Note that, since our A.P. surface flattening method is independent of the initial mapping, other surface parametrization methods can also be equally employed in order to achieve specific functionality relevant in applications. In
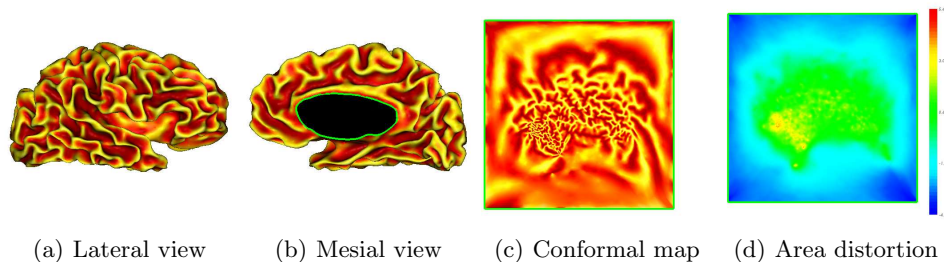

(a) Lateral view   (b) Mesial view   (c) Conformal map   (d) Area distortion

Fig. 1. Initial conformal flattening. (a) and (b) show the lateral and mesial views of a cortical hemisphere. (c) is the initial (conformal) flattening on a unit square. The area distortion is color-coded in (d).

the following, we describe the essential steps of the algorithm in the order they occur in the procedure.

3.1   Solving $\Delta g = h_2 - h_1$

For an area-preserving flattening, h_1 is always 1, whereas h_2 is the model area/parameter area ratio at play. When the total areas of the model and the domain are not equal, h_2 is subject to a normalization to make them equal. The result will therefore be of relative area preservation.

Given a continuous function on the surface, its discrete version is represented as a vector, defined on the vertex set V. To solve Eq. (8) on a triangular mesh, Δg is estimated at v_i in the same manner as [3]. Note that this discrete approximation requires acute angles. For obtuse ones, a proper remeshing procedure should be employed. As such, Eq. (8) can be written as a linear system: $Lx = b$, where $x = \boldsymbol{g}$, $b = \boldsymbol{h_2} - \boldsymbol{h_1}$, L represents the coefficient matrix of the discrete Laplace-Beltrami operator. As L is sparse, Eq. (8) can be solved efficiently in linear time. Fig. 2(a) shows the solved function g.
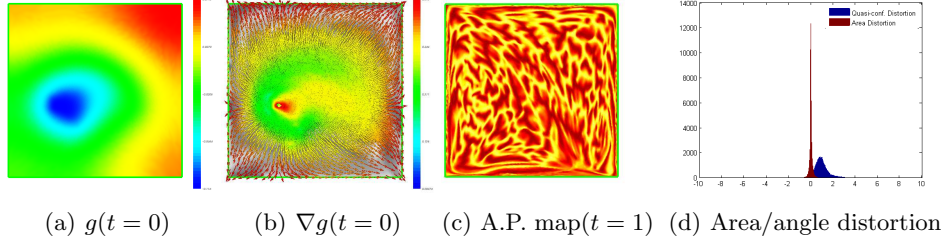
(a) $g(t=0)$ (b) $\nabla g(t=0)$ (c) A.P. map($t=1$) (d) Area/angle distortion

Fig. 2. Area-preserving flattening on a unit square. The solved function g is shown in (a) and its gradient vector field in (b). After time integration of a dynamic diffeomorphic flow following the vector field, the proposed A.P. flattening is obtained in (c). (d) shows the histograms of area and angle distortion metrics.

3.2   Computing ∇g

Now that the g is obtained, we can proceed to compute the corresponding gradient vector field on the triangulated domain. We consider a face f_{ijk} with its three corners lying at $\mathbf{v}_i$, $\mathbf{v}_j$, $\mathbf{v}_k$ in $\mathbf{R}^3$. Also, let $\mathbf{n}$ be a unit normal vector perpendicular to the plane spanned by f_{ijk}. By assuming linear interpolation within each triangle, the gradient vector can be easily computed by solving the 3×3 linear system:

$$\begin{bmatrix} \mathbf{v}_j - \mathbf{v}_i \\ \mathbf{v}_k - \mathbf{v}_j \\ \mathbf{n} \end{bmatrix} \nabla g = \begin{bmatrix} g_j - g_i \\ g_k - g_j \\ 0 \end{bmatrix}, \tag{10}$$

for which a closed-form solution exists. To obtain a unique vector at each vertex, ∇g at vertex v_i is defined as an average of the gradients of the adjacent faces, weighted by the incident angle α_{jk}^i of each face f_{ijk} at v_i. The resulting vector field is shown in Fig. 2(b).

3.3   Time Integration of $V(t)$

Recall that the Lie derivative is defined as the instantaneous change of forms evaluated at $\phi_t(x)$, which is a dynamic definition (See Eq. (4)). In fact, $V(t)$ only coincides with the flow vector field at $t = 0$. In other cases, $V(t)$ needs be transported accordingly. By an simple analytical integration, we have

$$\int_0^1 V(t)dt = \frac{\ln h_2 - \ln h_1}{h_2 - h_1} \nabla g. \tag{11}$$

When $h_1 = 1$, we have $\lim_{h_2 \to 1} \frac{\ln h_2 - \ln h_1}{h_2 - h_1} = 1$, which means that, when h_2 is sufficiently close to 1, the displacement vector field can be properly approximated by ∇g. Thus, the area-correcting process is divided into K sequential steps. With each step, the area element is only modified by an small amount δh towards the target setting such that the overall area adjustment is equal to $K\delta h$. More specifically, we let $h_2 = 1 + \delta h$ and $h_1 = 1$ as the input of the analytical integration (Eq. (11)), and the result gives the corresponding displacement for the area change of δh. This procedure is discretized into 50 steps for the results

shown. By choosing step length carefully, we can guarantee that the time integration converges. The final A.P. flattening of the cortical hemisphere is shown in Fig. 2(c). Fig. 3(d) presents the statistics of the area (darkred) and angle (blue) distortions; the employed metrics will be explained in Section 4. All boundary vertices can either be fixed, or evolved by the covariant derivative of $V(t)$ along ∂D. Empirically, the latter one gives better approximation to the continuous case, but less robust to degenerate mesh triangles.

3.4   Genus Zero Surface Flattening

The entire brain surface is often modeled as a topological sphere, i.e., a closed surface of genus zero, thus is preferred to be flattened on a unit sphere ($\mathbf{S}^2$) without any topological changes. Our method can easily adapt to this case as well. We first compute its initial mapping using the method described in [5]. The Lie advection flowing to the area preservation is then performed in the tangent spaces of the spherical domain. In practice, whenever a vertex moves out of the unit sphere, it is pulled back by $\tilde{v}_i = v_i/|v_i|$. Fig. 3(a) shows the entire cortex, while its spherical conformal mapping is shown in 3(b). Its A.P. flattening is shown in Fig. 3(c) and the corresponding histograms about the area (darkred) and angle (blue) distortions is presented in 3(d).
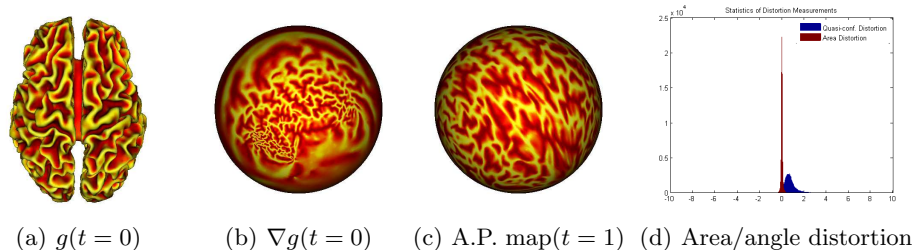

(a) $g(t = 0)$ (b) $\nabla g(t = 0)$ (c) A.P. map($t = 1$)   (d) Area/angle distortion

Fig. 3. Area preserving mapping of closed genus zero surface. (a) is the original brain surface model. Its spherical conformal mapping and subsequent A.P. flattening are shown in (b) and (c), respectively. Similarly, the statistics of the result is shown in (d).

4   Results

All brain surface models are reconstructed on the gray matter/white matter interface of 3D MRI brain volumes. To verify that the area elements (associated with the vertice) are preserved globally independent of the triangulation, we build up our distortion measures on their dual cells–triangular faces. To be specific, we examine both the area distortion and the quasi-conformal distortion per face over the mesh. The area distortion metric $\varUpsilon$ and the quasi-conformal distortion metric $\varLambda$ are computed respectively as follows:

$$\varUpsilon = \ln(\gamma_{max} \cdot \gamma_{min}), \quad \varLambda = \ln \frac{\gamma_{max}}{\gamma_{min}}, \tag{12}$$

where $(\gamma_{max}, \gamma_{min})$ are the larger and smaller eigenvalues of the Jacobian of the affine transformation that maps the domain triangle to the surface, normalized in such a way that the total area of the surface equals that of the domain. In both cases, a value of 0 indicates no distortion at all, while distortions can deviate on both sides.

In Fig. 2(d) and Fig. 3(d), the statistics of the area distortion (darkred) and the quasi-conformal ones (blue) are illustrated for both examples, respectively. From the distribution of metric Υ, we can see that the areas of triangles are clearly well preserved. Because for general manifolds, no isometric mapping exists except for a few special cases, quasi-conformal distortion is expected for A.P. maps, as indicated by the distribution of Λ.

In terms of runtimes, the underlying mesh of the cortical hemisphere model is composed of 48,287 faces and the A.P. flattening was obtained in 57 sec; the model of the entire cortex is with 99,736 faces and it took 122 sec. All experiments were conducted on an Intel T6600 2.20GHz laptop with 3GB RAM. Generally, the cost linearly depends on the size of the mesh and the number of discretized steps needed for the desired accuracy.

For a visual analytical framework that is designed to integrate various complementing neuroimaging modalities from multiple sources and for quantitative analyses, A.P. flattening exhibits unique advantages. In Fig. 4, the brain surface is sliced open along the medial plane at the bottom, without passing any significant anatomical features. The entire brain surface is then mapped to a rectangle with a height/width ratio of 1:2. To capture the statistics of abnormal brain activity, a sufficient number of sampling points should be grouped at certain resolution. Each unit forms a cortical element. Since the surface of the brain is uneven and varies across human subjects, subdividing it into a set of equal geometric elements is nontrivial. With A.P. flattening, we can easily define finite homotopic cortical elements on the parametric domain with simple isotropic grid. Comparing to other "quasi" area-preserving/conformal mapping, each element based on an A.P. flattening accounts for an identical amount of portion in the original brain surface. With these well-defined equiareal cortical elements, a variety of functional patterns can be readily quantified on a per element basis.
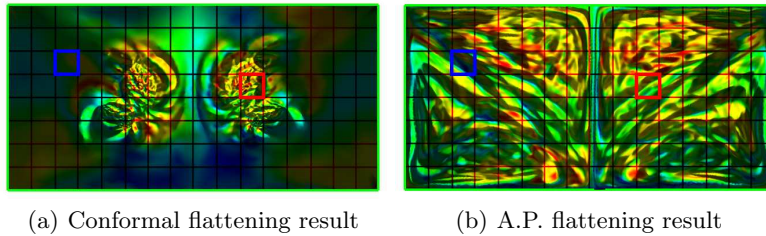

(a) Conformal flattening result   (b) A.P. flattening result

Fig. 4. Application of quantitative surface-based analytics. The color visualizes the integrated PET data. (a) shows the flattening result via conformal mapping, where most anatomical and functional meaningful areas are squeezed in a few square elements. In contrast, the cortical surface is more evenly sampled under A.P. flattening, as shown in (b).

5   Conclusion

In this paper, we have presented a special surface flattening methodology that is strictly area-preserving. Given an arbitrary initial parametrization, an A.P. result can be efficiently and uniquely obtained via the Lie advection of area forms along the domain. As our method only depends on the intrinsic surface geometry, it is insensitive to different discretized representations, including variations in triangulation and/or resolution. Our implementation strives to preserve the analytical ingredients of the computation as far as possible. As a result, our method is highly efficient and stable in getting consistent results, as opposed to an optimization approach. While conformal maps are known to be powerful for shape analysis, in a number of other scenarios, such as studies of neuronal density, cortical functional activation extent, and cortical thickness variations, where accurate modality sampling and statistical sensitivity are critical, areal preservation is highly preferable. Extensive comparison between our method and existing approaches in clinical practice will be conducted in the near future.

References

1. Angenent, S., Haker, S., Tannenbaum, A., Kikinis, R.: On area preserving mappings of minimal distortion. In: Djaferis, T., Schick, I. (eds.) System Theory: Modeling, Analysis, and Control. Kluwer, Holland (1999)
2. Angenent, S., Haker, S., Tannenbaum, A., Kikinis, R.: On the laplace-beltrami operator and brain surface flattening. IEEE Trans. Med. Imag. 8, 700–711 (1999)
3. Desbrun, M., Meyer, M., Alliez, P.: Intrinsic parameterizations of surface meshes. Comput. Graph. Forum 21, 209–218 (2002)
4. Fischl, B., Sereno, M., Dale, A.: Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. NeuroImage 9(2), 195–207 (1999)
5. Gu, X., Wang, Y., Chan, T., Thompson, P., Yau, S.: Genus zero surface conformal mapping and its application to brain surface mapping. IEEE Trans. Med. Imag. 23, 949–958 (2003)
6. Hurdal, M., Stephenson, K.: Discrete conformal methods for cortical brain flattening. NeuroImage 45, S86–S98 (2009)
7. Lévy, B., Petitjean, S., Ray, N., Maillot, J.: Least squares conformal maps for automatic texture atlas generation. ACM Trans. Graph. 21(3), 362–371 (2002)
8. Pons, J., Keriven, R., Faugeras, O.: Area preserving cortex unfolding. In: Proc. MICCAI 2004. pp. 376–383 (2004)
9. Zhu, L., Haker, S., Tannenbaum, A.: Flattening maps for the visualization of multi-branched vessels. IEEE Trans. Med. Imaging 24(2), 191–198 (2005)