

LMap: Shape-Preserving Local Mappings for Biomedical Visualization

Saad Nadeem, Xianfeng Gu, and Arie Kaufman, *Fellow, IEEE*

Abstract—Visualization of medical organs and biological structures is a challenging task because of their complex geometry and the resultant occlusions. Global spherical and planar mapping techniques simplify the complex geometry and resolve the occlusions to aid in visualization. However, while resolving the occlusions these techniques do not preserve the geometric context, making them less suitable for mission-critical biomedical visualization tasks. In this paper, we present a shape-preserving local mapping technique for resolving occlusions locally while preserving the overall geometric context. More specifically, we present a novel visualization algorithm, LMap, for conformally parameterizing and deforming a selected local region-of-interest (ROI) on an arbitrary surface. The resultant shape-preserving local mappings help to visualize complex surfaces while preserving the overall geometric context. The algorithm is based on the robust and efficient extrinsic Ricci flow technique, and uses the dynamic Ricci flow algorithm to guarantee the existence of a local map for a selected ROI on an arbitrary surface. We show the effectiveness and efficacy of our method in three challenging use cases: (1) multimodal brain visualization, (2) optimal coverage of virtual colonoscopy centerline flythrough, and (3) molecular surface visualization.

Index Terms—Biomedical visualization, virtual colonoscopy, multimodal brain visualization, molecular surface visualization, shape-preserving mapping

1 INTRODUCTION

Visualization is critical for diagnosing various medical conditions (e.g., different types of cancer) through non-invasive imaging. However, 3D visualization of medical organs and various biological structures is challenging because of their complex 3D geometry and the resultant occlusions. For example, deep folds in the brain are occluded by the gyri (ridges) on the brain cortical surface (Figure 1a). Global spherical and planar mapping techniques help to resolve these occlusions by mapping the complex 3D geometry to a canonical representation such as a sphere or a plane, respectively, with minimal angle and/or area distortions. The removal of both occlusions and the geometric context makes it difficult to correlate the features between the original and the mapped surfaces (Figure 1b), making these techniques less useful for mission-critical biomedical visualization tasks.

In effect, spherical and planar mesh mapping techniques are mostly restricted or relegated to segmentation, registration, and classification tasks. Moreover, these techniques assume a certain topology for the input surface, for example, spherical mapping techniques mandate a genus 0 closed surface input, whereas planar mapping techniques require a genus 0 surface with a boundary. These restrictions force the users to conform the data to the assumed input requirements via pre-processing (e.g., hole-filling, handle removal, feature extraction, etc.) to make these techniques work. For example, Krone et al. [18] had to remove important channel information from the molecular surfaces for their spherical mapping technique to work, and Zeng et al. [35] had to perform topological denoising (removal of fake handles) on the colon surface and extract consistent teniae coli (often not possible

due to collapsed regions or segmentation artifacts [21]) for their flattening approach. Similar problems underline the multimodal domain where the information from all the modalities is computed in the native space and transformation of the original geometry (to a sphere or a plane) requires re-computation of all the modalities in the mapped space [20].

Ideally, a local mapping algorithm is required to deal with the shortcomings of the global mapping techniques. We prescribe the following desirable properties for our local mapping algorithm (modified from the ones outlined in Rocha et al. [26] for local parameterization): (1) it should be independent of the surface parameterization since many scientific datasets are not represented parameterically; (2) it should not rely on any global information from the underlying mesh since this is expensive to compute and depends on the mesh type; (3) it should be efficient since we can require a large number of parameterizations and deformations of several regions-of-interest (ROIs) on the surface; (4) it should be angle-preserving (conformal) for the selected ROI and should not affect the rest of the surface; and (5) it should work for arbitrary surfaces (open/closed, smooth/convoluted, arbitrary-genus).

In this work, we present a technique, LMap, to parameterize and conformally deform local ROIs on arbitrary surfaces for occlusion-free visualization while preserving the overall geometric context (Figure 1d). More specifically, for a selected ROI on a mesh, LMap uses a novel extrinsic Ricci flow algorithm for parameterizing and conformally deforming this selected ROI. The presented extrinsic Ricci flow algorithm entails the following steps: (1) compute local Gaussian curvature on the selected ROI, (2) obtain new edge lengths for the computed Gaussian curvature using dynamic Ricci flow [22], and (3) deform the position of each ROI vertex on the mesh, according to the new edge lengths, to obtain the conformally flattened ROI on the given mesh. The resultant output mesh resolves occlusions for selected ROIs while keeping the rest of the surface intact and untouched. Two

• Saad Nadeem, Xianfeng Gu and Arie Kaufman are with the Department of Computer Science, Stony Brook University, Stony Brook, NY 11794-2424. E-mail: {sanadeem, gu, ari}@cs.stonybrook.edu.




Fig. 1: Spherical and local mappings for brain cortical surface overlaid with mean curvature. (a) Front and (e) back of the brain cortical surface overlaid with mean curvature. (b) Front and (f) back of the angle-area preserving spherical mapping [22]. (c) Front and (g) back selected regions-of-interest, colored in green. (d) Front and (h) back of the local mapping result from our algorithm. The dark blue lines in locally mapped surfaces (d) and (h) correspond to the sulcal fundi which have been shown to be robust and reproducible cortical landmarks for registration and group-level analysis. LMap computation for the selected ROIs (78952 vertices for the front ROI and 82291 vertices for the back ROI) took 2.7 seconds.

important implications of the presented LMap algorithm are as follows: (1) the use of dynamic Ricci flow guarantees the existence of a stable (robust to perturbations in the selected ROIs) local map for a selected ROI on an arbitrary surface, and (2) the use of Gaussian curvature for the local ROI results in a stable local map, as opposed to the use of mean curvature for local ROIs which can lead to self-intersections (as shown in Figure 2h).

In essence, LMap does not require any global information from the underlying mesh to deform the selected ROIs, and LMap is not dependent on the underlying topology or the type of the mesh. Moreover, unlike recent approaches which can only parameterize sphere-masked smooth local regions on arbitrary surfaces [26], LMap is applicable for smooth as well as convoluted local (sharp bend) ROIs on arbitrary surfaces. It also does not require any prior hole-filling or topological denoising of the input surface. Intuitively, LMap acts like a magic lens for selected ROIs on arbitrary surfaces, where it can be used to locally parameterize, and overlay textures and different types of information from multimodal and multivariate data. LMap also allows for local shape-preserving deformations for better visualization of complex ROIs.

We show the efficacy and effectiveness of LMap in three challenging use cases: (1) multimodal brain visualization, (2) optimal coverage in virtual colonoscopy centerline flythroughs, and (3) molecular surface visualization. In the context of multimodal brain visualization, neuroscientists want to study specific cortical structures (which are highly convoluted and difficult to visualize in their original form) and the connectivity patterns between these structures. Unlike previous spherical parameterization methods [20], [22] which mandate the transformation of the all the modalities to the underlying spherical space, LMap allows the simplified visualization of these specific structures and the connectivity patterns in the native space while preserving the overall geometric context. For virtual colonoscopy use case, radiologists want to

localize and study the polyps (precursors of colon cancer) as accurately as possible. Current angle-preserving colon flattening approaches [19], [21], [35] induce considerable area distortion, require teniae coli extraction (for cutting) and distort the geometry significantly, making the polyp localization difficult. In contrast, LMap can be used to only deform the detected haustral folds [21] in place while leaving the original geometry of the polyps and the overall colon intact, thus allowing for accurate polyp localization and measurements. Finally, in the context of molecular surface visualization, different cavities on the boundary and in the interior of the molecules need to be visualized as potential binding sites for other molecules. Current simplification approaches [18], in this context, use spherical parameterization to visualize the individual boundary cavities at the expense of removing the critical interior cavities and filling the resultant holes which induces considerable angle and area distortion during the parameterization; higher genus surfaces can induce considerably higher angle and area distortion due to higher loss in interior information. LMap, on the other hand, can be used for conformally deforming only the boundary cavities without removing the critical interior cavities and without affecting the rest of the molecular surface.

The main contributions of this paper are summarized as follows:

- We present a novel local mapping algorithm, LMap, for visualizing local regions-of-interest on arbitrary genus closed/open surfaces. LMap deforms local regions efficiently with minimal angular distortion using robust and theoretically-sound extrinsic Ricci flow. Moreover, the use of established dynamic discrete surface Ricci flow guarantees the existence of the local for the prescribed Gaussian curvature, and greatly improves the stability of the algorithm.
- We show the efficacy and effectiveness of LMap in three challenging use cases: (1) multimodal brain visualization,

(2) optimal coverage in virtual colonoscopy centerline flythroughs, and (3) molecular surface visualization. We also compare the LMap results with the state-of-the-art visualization algorithms in these domains.

2 RELATED WORK

Various global mapping techniques have been proposed in the past for medical visualizations, especially for colon, brain, and molecular surfaces. Most of these techniques require removal of segmentation artifacts, topological denoising (e.g., removal of fake handles), hole-filling (e.g., removal of channels mandates hole-filling in parameterized molecular surface visualization), and extraction of difficult-to-find features (e.g., consistent teniae coli extraction for colon flattening).

Colon flattening is a method in which the entire inner surface of the colon is displayed as a 2D image and thus all the occlusions are resolved but the geometric context is not preserved. The flattening requires the consistent teniae coli extraction which is often not possible and the folds in the resultant 2D image are cut at sharp bends [21], reducing the overall usefulness of this flattened result in the medical visualization context. Initial attempts to flatten the colon surface include iterative methods based on electrical field lines [31], cartographic [23], and cylindrical projections [3], [4]. However, most of these methods do not preserve the local shape. Conformal mapping, an approach where the local angles are preserved, has recently been used for colon flattening with promising results [11], [12], [14], [35]. The conformally flattened colon was used in the detection of colonic folds [14] and for supine-prone colon registration [35]. To overcome the need of topological denoising for conformal colon flattening, Gurijala et al. [11] used a heat diffusion metric and obtained promising results. A geometric context-preserving colon flattening approach was presented by Marino et al. [19] which was angle-preserving but induced large area distortion; however, the approach could not handle topological noise and required the prior consistent teniae coli extraction, mostly done manually. In our earlier work [21], we presented an automatic method to segment folds and extract consistent teniae coli for flattening, doing away with the tedious cutting process. However, segmentation artifacts remain a problem since it is difficult to find a consistent cutting loop through the collapsed regions of the colon. Moreover, current colon flattening approaches are only angle-preserving and hence, the induced area-distortion can be large, making them less useful for polyp visualization on the resultant 2D flattened images. LMap can be used to only deform the detected haustral folds [21] in place while leaving the original geometry of the polyps and the overall colon intact, thus allowing for accurate polyp localization and measurements.

Conformal mapping techniques have also been used for brain visualization, such as the circle packing-based method [15], finite element method [2], [16], [29], and spherical harmonic map method [6], [10], [32]. Zhao et al. [36] presented an angle-area preserving planar brain flattening approach using optimal mass transport. In our earlier work [22], we developed an angle-area preserving 3D spherical mapping technique for the brain and other genus 0 surfaces. The problem with these approaches is that because of the loss in geometric context, they are primarily used for registration, segmentation and classification tasks but not for brain visualization. In contrast, LMap allows simplified

visualization of the local regions of interest while preserving the overall geometric context

In computational chemistry, molecules are often visually compared based on local surface attributes. However, to the best of our knowledge all the previous work for making this local comparison rely on global mapping techniques, which do not preserve the overall geometric context, and hence can make it difficult to make a finer comparison in the case of large regions with uniform attribute values. Previous works leveraged spherical mapping via deformable models [24], spherical coordinates [25], conformal mapping [13], surface projections [17], and map projections [18]. Krone et al. [18] presented a hole-filling technique to reduce the genus surface to genus 0 and map the resultant genus 0 surface to a sphere for visualization. Unfortunately, the simplification to genus 0 surface was done at the expense of removing critical channel information and filling the resultant holes, which induces significantly high angle and area distortion; higher genus surfaces can induce considerably higher angle and area distortion due to higher loss in interior information. LMap, on the other hand, can be used for conformally deforming only the boundary cavities without removing the critical interior cavities and without affecting the rest of the molecular surface.

The method closest to our LMap algorithm is spin transformations [7]. Spin transformations require global information, i.e. mean curvature computation for the entire mesh, to allow for conformal deformations for selected ROIs on this mesh; restricting the mean curvature computation to a local region while ignoring the rest of the surface leads to self-intersections (as shown in Figure 2h). In contrast, LMap only requires local Gaussian curvature information to parameterize and conformally deform selected ROIs, and does not affect the rest of the surface, hence making it more useful for biomedical visualization tasks than the previous methods. Simple smoothing constrained to the boundary of the selected ROI does not preserve shape and hence, is not useful in the present context. LMap is inspired from the recent works in conformal parameterization domain ([1], [28]).

3 COMPUTATIONAL ALGORITHM

This section introduces the computational algorithm for the intrinsic and extrinsic Ricci flow. The intrinsic Ricci flow [22] is used to obtain the new lengths for the target Gaussian curvature and the extrinsic Ricci flow is used to compute the corresponding deformed mesh. More details on the intrinsic Ricci flow can be found in our previous paper [22].

3.1 Intrinsic Surface Ricci Flow

In practice, surfaces are represented as triangular meshes. A *discrete metric* on a mesh is the edge length function, denoted as $l : E \rightarrow \mathbb{R}^+$, which satisfies the triangle inequality. The *discrete Gauss curvature* is the angle deficit, defined on vertices, $K : V \rightarrow \mathbb{R}$,

$$K(v) = \begin{cases} 2\pi - \sum_{jk} \theta_i^{jk}, & v \notin \partial M \\ \pi - \sum_{jk} \theta_i^{jk}, & v \in \partial M \end{cases}, \quad (1)$$

where θ_i^{jk} is the corner angle at v_i in the face $[v_i, v_j, v_k]$, and ∂M represents the boundary of the mesh.

The discrete Gaussian curvature is determined by the discrete Riemannian metric via the cosine law,

$$l_i^2 = l_j^2 + l_k^2 - 2l_j l_k \cos \theta_i. \quad (2)$$

According to the Gauss-Bonnet theorem, the total Gaussian curvature is a topological invariant and equals the product of 2π and the Euler characteristic number χ ,

$$\sum_{v \notin \partial \Sigma} K(v) + \sum_{v \in \partial \Sigma} K(v) = 2\pi\chi(\Sigma). \quad (3)$$

The *cotangent edge weight* plays an important role. Given an interior edge $[v_i, v_j]$ adjacent to two faces $[v_i, v_j, v_k]$ and $[v_j, v_i, v_l]$, the cotangent weight is defined as

$$w_{ij} = \cot \theta_k^{ij} + \cot \theta_l^{ji}. \quad (4)$$

If the edge is on the boundary, adjacent to the face $[v_i, v_j, v_k]$, then the cotangent weight is

$$w_{ij} = \cot \theta_k^{ij}. \quad (5)$$

A triangulation of the mesh is called *Delaunay* if all cotangent edge weights are non-negative.

Given a triangular mesh M , the *discrete conformal factor* is a function defined on each vertex $\mathbf{u} : V \rightarrow \mathbb{R}$, the length of an edge $[v_i, v_j]$ is given by

$$l_{ij} = e^{u_i} \beta_{ij} e^{u_j}, \quad (6)$$

where β_{ij} is the initial edge length.

Discrete Surface Ricci Flow: The discrete surface Ricci flow is defined as

$$\frac{du_i(t)}{dt} = \bar{K}_i - K_i(t), \quad (7)$$

where $\bar{K}_i$ is the target curvature at the vertex v_i , and the discrete metric is given by Eqn. (6).

Furthermore, the discrete Ricci flow is the negative gradient flow of the discrete Ricci energy:

$$E_\Sigma(\mathbf{u}) = \int^\mathbf{u} \sum_{i=1}^n (\bar{K}_i - K_i) du_i. \quad (8)$$

The gradient of the Ricci energy is $(\bar{K}_i - K_i)^T$, the Hessian matrix consists of cotangent edge weights, given in Eqn. (4),

$$\frac{\partial K_i}{\partial u_j} = \frac{\partial K_j}{\partial u_i} = w_{ij}, \quad (9)$$

and the diagonal elements are

$$\frac{\partial K_i}{\partial u_i} = -\sum_j w_{ij}. \quad (10)$$

Dynamic Surface Ricci Flow [22]: Given a target curvature $\bar{K} : V \rightarrow \mathbb{R}$, satisfying the Gauss-Bonnet condition (3), conventional Ricci flow does not guarantee the existence of the solution. During the conventional discrete surface Ricci flow, some triangles may become degenerated; namely the triangle inequality does not hold on some faces, hence the flow terminates. In order to overcome the instability of conventional surface Ricci flow, we use dynamic surface Ricci flow [22]: during the Ricci flow, the triangulation of the mesh is updated to be Delaunay with respect to the current Riemannian metric.

The Delaunay triangulation can be obtained by a finite number of edge flipping operations: for each edge, we flatten the two adjacent faces to form a planar quadrilateral, then swap the edge to the shorter diagonal. During the optimization, the triangulation is updated to be Delaunay at each step; details can be found in Algorithm 1. Preserving Delaunay triangulation is crucial to ensure the existence of the solution to the dynamic Ricci flow. Given a target curvature $\bar{K}$ satisfying the Gauss-Bonnet condition

in Eqn. (3), and for each vertex $\bar{K}_i \in (-\infty, 2\pi)$, there exists a solution to the dynamic Ricci flow [9]. The solution is unique up to a constant.

Algorithm 1: Intrinsic Surface Ricci Flow

Require: The input mesh M , target curvature $\bar{K}$, threshold ε
Ensure: The edge length which realizes the target curvature
 Compute the initial edge lengths $\{\beta_{ij}\}$;
 Initialize the conformal factor to be zeros;
while true do
 Compute the edge lengths using Eqn. (6);
 Update the triangulation to be Delaunay by edge flipping;
 Compute the corner angles using Eqn. (2);
 Compute the cotangent edge weights using Eqns. (4) and (5);
 Compute the vertex curvature using Eqn. (1);
 if $\forall |\bar{K}_i - K_i(\mathbf{h})| < \varepsilon$ **then**
 Break;
 end if
 Compute the gradient of the Ricci flow using Eqn. (7);
 Compute the Hessian of the Ricci energy using Eqns. (9) and (10);
 Solve the linear system

$$Hess(\mathbf{u})\delta\mathbf{u} = \nabla E(\mathbf{u})$$

$$\mathbf{u} \leftarrow \mathbf{u} + \delta\mathbf{u};$$

end while
return the edge length $\{l_{ij}\}$

3.2 Extrinsic Surface Ricci Flow

Given a triangle mesh $M = (V, E, F)$, we select a region of interest (ROI) $\Omega \subset M$. We choose the step length as $1/n$, where n is a positive integer. At the k -th step, we set the target curvatures for vertices inside Ω ,

$$\bar{K}(v_i) = \left(1 - \frac{k}{n}\right) K(v_i), \quad \forall v_i \in \Omega, \quad (11)$$

where $K(v_i)$ is the initial discrete Gaussian curvature. We use the intrinsic discrete surface Ricci flow algorithm (Algorithm 1) to compute the corresponding edge length, denoted as $l^k : E \rightarrow \mathbb{R}_{>0}$. Then we adjust the positions of the vertices to satisfy the edge length constraints.

At the k -th step, the mesh is M^k , the position for the vertex v_i is p_i^k , and the deformation is along the normal field of the surface. Let $\lambda^k : V \rightarrow \mathbb{R}$ be a function defined on the vertex set, the deformation of v_i is represented as

$$p_i^{k+1} \leftarrow p_i^k + \lambda_i^k n_i^k, \quad (12)$$

where n_i^k is the normal to the mesh M^k at the vertex v_i . The area of the face $[v_i, v_j, v_k]$ is given by

$$s_{ijl}^k = \frac{1}{2} |(p_j^k - p_i^k) \times (p_l^k - p_i^k)|.$$

The normal to the face can be computed as

$$n_{ijl}^k = \frac{1}{2s_{ijl}^k} (p_j^k - p_i^k) \times (p_l^k - p_i^k).$$




Fig. 2: Spin transformations comparison for the molecular surface in Figure 9a. (a) Mean curvature and *global* spin transformations with varying scale factors (b) 2, (c) 6, and (d) 10. (e) Mean curvature is set to zero everywhere except the selected ROI and corresponding *local* spin transformations with varying scale factors (f) 2, (g) 6, and (h) 10. (h) Surface self-intersects (black bounding box). (i) Selected ROI and the corresponding LMap deformations (using Gaussian curvature) after (j) 1, (k) 2, and (l) 5 steps.

The normal to the vertex v_i is given by

$$d_i^k = \sum_{jl} s_{ijl}^k n_{ijl}^k, \quad n_i^k = \frac{d_i^k}{|d_i^k|}. \quad (13)$$

The edge length of an edge $[v_i, v_j]$ is l_{ij}^k , the square of the length is L_{ij}^k , then the energy for the optimization is given by

$$E(\lambda^k) := \sum_{[v_i, v_j] \in E} \left[\langle p_i^{k+1} - p_j^{k+1}, p_i^{k+1} - p_j^{k+1} \rangle - L_{ij}^k \right]^2. \quad (14)$$

By direct computation, the gradient of the energy can be calculated as

$$\frac{\partial E(\lambda^k)}{\partial \lambda_i^k} = \sum_{v_i \sim v_j} (L_{ij}^{k+1} - L_{ij}^k) \langle n_i^k, p_i^k - p_j^k \rangle \quad (15)$$

The algorithmic details can be found in Algorithm 2. The vertices outside the selected ROI are fixed in place and hence, are not affected by the local mapping of the given ROI.

3.3 Evaluation Method

The LMap algorithm was developed using generic C++ on the Windows 7 platform. The Matlab C++ library was used for solving the linear system. All of the experiments are conducted on a workstation with a Core 2 Quad 2.50 GHz CPU with 8GB RAM.

We compare LMap against the spin transformations technique [7]. We refer to the spin transformations with mean curvature for the whole input surface as *global* spin transformations and the

Algorithm 2: Extrinsic Surface Ricci Flow

Require: The input mesh M and the ROI $\Omega \subset M$

Ensure: The deformed mesh M such that Ω is flat

$k \leftarrow 0$

while $k < n$ **do**

Set the target curvature for vertex in Ω using Eqn. (11)

Use Alg. 1 to compute the edge lengths l^k

Minimize the energy in Eqn. (14) using gradient descent method, using the formula Eqn. (15)

Deform the position of each vertex in Ω using Eqn. (12)

$k \leftarrow k + 1$;

end while

return the deformed mesh

ones with mean curvature restricted to the selected ROI as *local* spin transformations. For a selected ROI on the molecular surface shown in Figure 9b, we compute the corresponding local spin transformations where mean curvature is set to zero everywhere except the selected ROI (Figures 2e–2h) and global spin transformations (Figures 2a–2d) with varying scale factors. We can observe that the local spin transformations for a selected ROI can significantly distort the scale of the rest of the surface and can lead to self-intersections with larger scale factors. In contrast, LMap does not affect the rest of the surface while parameterizing and deforming the selected ROI, as shown in Figures 2i–2l. Moreover,




Fig. 3: Stability of LMap with ROI perturbations. Selected ROIs, (a) and (c), and their corresponding local mappings, (b) and (d).




Fig. 4: Sharp Bends. (a) Selected ROIs and the (b) corresponding local mappings with normal mapped shading from (a).

the LMap is stable to perturbations of ROI, as shown in Figure 3. LMap can also handle sharp bends, as demonstrated in Figure 4.

In theory, the Ricci flow method induces a conformal mapping. A conformal map preserves angles and local shapes. The shape-preserving property is crucial for visualization purposes. In order to evaluate the conformality of the mapping result, we compute both angle and area distortions. The area distortion is computed as follows. Assume the parameterization is $\phi : M \rightarrow \mathbb{S}$. For each vertex v_i , the *area distortion* is defined as

$$\epsilon_i := \log \frac{\sum_{j,k} A([\phi(v_i), \phi(v_j), \phi(v_k)])}{\sum_{j,k} A([v_i, v_j, v_k])}$$

where $A(\cdot)$ represents the area of a triangle, and $[v_i, v_j, v_k]$ is the triangle formed by v_i, v_j, v_k . We then plot the histograms of ϵ_i . Similarly the *angle distortion* at a corner angle is given by

$$\eta_{ijk} := \log \frac{\angle \phi(v_i) \phi(v_j) \phi(v_k)}{\angle v_i v_j v_k},$$

we then plot the histograms of η_{ijk} .

The angle-preserving (conformal) mapping should ideally be close to zero angle distortions everywhere, whereas the area-preserving mapping should be close to zero area distortions everywhere. We compute the area and angle distortion histograms for the local selected ROI (Figure 9b) on mapped molecular surfaces computed using LMap (Figure 9b), global spin transformation




Fig. 5: Area and angle distortion histograms. (a) Angle and (b) area distortion histograms for LMap on the selected ROI in Figure 2l. (c) Angle and (d) area distortion histograms for local spin transformation on the same ROI with scale factor 6, as shown in Figure 2g. (e) Angle and (f) area distortion histograms for global spin transformation on the same ROI with scale factor 6, as shown in Figure 2c.

with scale factor 6 (Figure 2c) and local spin transformation with scale factor 6 (Figure 2g). As shown in Figure 5, for a given ROI, LMap preserves angle and area better than both global and local spin transformations.

4 CASE STUDIES

We present the following three challenging case studies for our LMap algorithm.

4.1 Multimodal Brain Visualization

The human brain is complex in its shape and connectivity, even at scales at which magnetic resonance imaging can measure. Visualization is used for exploring possible relationships between various brain structures across scan types (multimodal data) or statistical variables (multivariate data). However, the convoluted




Fig. 6: Multimodal brain visualizations using local mappings. (a) Brain cortical surface overlaid with cortical parcellation. The same color regions on the left and right hemispheres represent the counterparts on the respective hemispheres. (b) Structural connectivity from the diffusion MRI data is represented as fiber tracts, seeded at the corpus callosum. Left hemisphere (c) original surface and the locally mapped surface at the selected region (d) using standard surface shading and (e) using normal mapped shading from the original mesh. (f) The longest structural fiber connections to this selected region are shown in yellow with corresponding locally mapped left hemisphere surface. (g) The corresponding connections and the locally mapped surfaces on the right hemisphere. (h) Functional connectivity data is overlaid on the locally mapped surfaces, with white representing the most functionally-activated regions and black indicating the least; the seed region is the upper white region on the left locally mapped surface. LMap computation for the two ROIs took 1.7 secs.




Fig. 7: Multimodal brain visualizations comparison. (a) Angle- and area-preserving spherical parameterization [22] of brain cortical surface in Figure 6a; the red curve represents the conformally welded boundaries of the left and right hemispheres. (b) The longest structural fiber connections to the same selected regions, as shown in Figure 6h.

brain cortical structures induce considerable occlusions making it difficult to visualize these structures in their original form.

Current global brain mapping techniques resolve these occlusions [22] and allow for exploration of multiple modalities on a brain surface [20] but at the expense of losing the geometric context. Moreover, the biggest problem with the global brain mapping techniques is that all the modalities have to be mapped to the common simplified representation and this can be a major bottleneck. An example of the multimodal brain visualization [20] based on angle- and area-preserving spherical parameterization is presented in Figure 7.

In contrast, LMap overcomes these limitations by locally mapping selected ROIs while preserving the overall geometric context and allowing for simultaneous visualization of cortical features, along with structural and functional connectivity information from diffusion and functional magnetic resonance imaging (MRI) data respectively in the native space, as shown in Figure 6. We used the Human Connectome Project (HCP) data [30] along with DSI Studio [34] to obtain the highest fidelity structural, functional, and diffusion data for our visualizations.

We also consulted a neuroscientist for feedback on our multimodal brain visualizations. Since the neuroscientists want to study specific cortical structures and the connectivity patterns




Fig. 8: Optimal coverage in virtual colonoscopy. (a) Virtual colonoscopy mesh with two polyps, one in view and another one occluded by a far fold. (b) Highlighted folds between the near and the far polyps are conformally deformed via LMap and (c) the far polyp is exposed (highlighted with a yellow bounding box) while preserving the overall geometric context. The coverage of this colon segment increased from 72% in (a) and (b) to 94% using LMap in (c). (d) Polyp occluded by the far fold and the corresponding LMap on the (e) highlighted folds in green remove the occlusion and (f) exposes the polyp (highlighted with a yellow bounding box). The coverage of this colon segment is increased from 76% in (d) and (e) to 91% in (f). (g) Polyp on the fold. (h) Highlighted polyp in green and the corresponding fold are conformally deformed via LMap (and pushed towards the colon wall) to increase coverage and the result is shown (i) using standard surface shading and (j) using normal mapped shading from the original mesh. The shape of the polyp is preserved in (j), owing to the angle-preserving property of LMap. The folds in (b), (e), and (h) are automatically detected using the method presented in our earlier work [21]. (k) Angle-preserving flattening [21] of colon segment in (f) with the polyp highlighted with a yellow bounding box. (l) Flattened colon segment in (g) with the polyp highlighted with a yellow bounding box. LMap computation for ROIs in (b), (e), and (h) took 1.8, 2.3, and 1.2 secs, respectively.

between them, the ability of LMap to focus on the simplification of these specific structures while preserving the overall geometric context was highly appreciated. Moreover, the local mappings for arbitrary-sized ROIs in the native space were also acknowledged, along with the possibility of visualizing subcortical features in the native space, and accurately overlaying the multimodal information from the resting-state and task-based functional MRI data, and

multivariate information from geometric measures (travel depth, curvature, area, thickness) and important landmarks such as sulcal pits and fundi (Figure 1). Currently, neuroscientists visualize deep folds on brain hemispheres by using inflated cortical surfaces or via planar transformation by removing the software-added tiling (shown in Figure 1e) and manually cutting along the cortical surface to lay out the brain on a plane in the original superior




Fig. 9: Molecular surface local mappings. (a) A small protein (genus 0) colored by temperature factor (PDB ID: 1RWE), with red the highest and blue the lowest temperature factors. Local mapping for a selected ROI (b) using standard surface shading (c) and using normal mapped shading from the original mesh (d). (e) Protein (genus 1) with a channel (PDB ID: 2BT9). The coloring shows the three amino acid chains that form the protein. Local mapping for selected ROIs (f) using standard surface shading (g) and using normal mapped shading from the original mesh (h). (g) Protein (genus 65) with an intricate network of channels (PDB ID: 1GKI). The coloring shows the six amino acid chains that form the protein. Local mapping for selected ROIs (j) using standard surface shading (k) and using normal mapped shading from the original mesh (l). LMap computation for ROIs in (b), (f), and (j) took 1.3, 2.0, and 3.8 secs, respectively.

view (shown in Figure 1a). We can achieve the tedious planar transformation by just selecting a larger region without requiring any further input.

4.2 Optimal Coverage in Virtual Colonoscopy

Virtual colonoscopy (VC), or CT colonoscopy, is a non-invasive cancer screening technique to navigate and inspect a 3D reconstructed colon surface for polyps (precursors of colon cancer) and if found, to localize and accurately measure these polyps. As part of the standard VC protocol, a radiologist inspects the surface while flying through the colon along the centerline from rectum to cecum and back in both supine (face up) and prone (face down) patient colon datasets. These four centerline flythroughs provide the optimal coverage of the colon surface and hence, minimize the probability of missing a polyp occluded by a fold. In total, a radiologist can spend between 15-30 minutes on these four flythroughs. With the deluge of non-invasive imaging, this interpretation time is a major bottleneck for these radiologists.

Global planar parameterization (or flattening approaches) can be used to completely resolve the occlusions in this context [21] and maximize the coverage. However, the problem with these approaches is that the resultant surface completely loses the geometric context and induces significant areal distortion since all these approaches are only angle-preserving (Figures 8k and 8l), making these approaches less suitable for polyp localization and visualization tasks; it is for this reason that these flattening approaches are only used for registration purposes. Moreover, these approaches require teniae coli extraction (for boundary mapping) which might not be possible in most cases [21] where the colon is under-distended (a common occurrence).

We show the effectiveness of LMap in this context by deforming only the detected haustral folds [21] in place while leaving the original geometry of the polyps and the overall colon intact, thus allowing for accurate polyp localization and measurements, as shown in Figure 8. The coverage of the colon segment, shown in Figure 8a, increased from 72% to 94% using LMap on the detected folds, in Figure 8c. The coverage of the colon




Fig. 10: Molecular surface maps [18] comparison. (a) A protein (PDB ID: 1RWE) is mapped to (b) a sphere using the force-directed method and (c) projected using the Lambert equal-area projection (as outlined in Krone et al. [18]). (d) A protein (PDB ID: 2BT9), with channel removed and the resultant holes filled (shown in black), is mapped to (e) a sphere with the parameter-based method and (f) projected using the Plate Carrée projection (as done in Krone et al. [18]).




Fig. 11: Molecular surface local mappings for cavities (a) and (c), colored in green. The corresponding local mappings are shown in (b) and (d).

segment, shown in Figure 8d, is increased from 76% to 91% using LMap on the detected folds, in Figure 8f. The polyps on the colon wall are left intact (distortion-free), making the polyp localization, visualization and diagnosis in the resultant maps as accurate as possible. The polyps can be located on both the colon wall and the folds and hence, shape-preservation during the local mapping is critical to avoid loss of such important information. As demonstrated in Figure 8g, a fold is detected with a polyp located on it using our algorithm [21]. The detected fold and the polyp located on it are locally deformed using LMap, as shown in Figure 8h. Due to the angle-preservation and local deformation

properties of LMap, we can visualize the true shape of the locally mapped polyp with minimal possible area distortion, as shown in Figure 8i. The optimal coverage in a single flythrough has a potential to significantly reduce the interpretation time and the associated costs.

We conducted an interview with a VC expert to receive feedback on our current LMap colon visualizations. The expert found the current visualizations with geometric context preservation very promising, especially for training new radiologists on VC systems. The expert reiterated our concern about the current state-of-the-art colon flattening approaches being angle-preserving only which can induce a large area distortion on the flattened surface, making them less desirable for polyp diagnosis. Since LMap only maps local ROIs (folds in the case of VC) the area distortion is much smaller than that induced with a completely flattened colon.

We use real CT colon data from the publicly available National Institute of Biomedical Imaging and Bioengineering (NIBIB) Image and Clinical Data Repository provided by National Institute of Health (NIH). We perform electronic cleansing incorporating the partial volume effect [33], segmentation with topological simplification [14], and reconstruction of the colon surface as a triangular mesh via surface nets [8] on the original CT images in a pre-processing step. Though the size and resolution of each CT volume varies between clinical datasets, the general data size is approximately $512 \times 512 \times 450$ voxels with a resolution of approximately $0.7 \times 0.7 \times 1.0$ mm.

4.3 Molecular Surface Visualization

Solvent excluded molecular surfaces are among the most prevalent visual representations for analyzing molecules, especially in biochemistry and structural biology molecular simulations. Different surface models represent different properties of the molecule and additional information can be overlaid on the surface, for example, to show physico-chemical properties of the underlying atoms [18].

The cavities on the boundary and in the interior of the molecule represent important information about the potential binding sites for other molecules making these cavities critical to study via visualization.

We demonstrate the effectiveness of LMap on three molecular surface datasets from Protein Data Bank [5], courtesy Krone et al. [18]. The two genus 0 (Figure 9a) and genus 1 (Figure 9e) surface datasets were used by Krone et al. [18] but the third genus 65 (Figure 9i) surface dataset was not; Krone et al. [18] computed the solvent excluded triangular mesh surfaces using the MSMS software by Sanner et al. [27]. LMap results for the corresponding ROIs (Figures 9b, 9f and 9j) using standard surface shading and using normal mapped shading from the original mesh are shown in Figure 9. LMap does not require the removal of interior cavities. It locally parameterizes and deforms only selected ROIs (e.g., boundary cavities) for visualization with minimal angle and area distortion while preserving the overall geometric context.

In contrast, current state-of-the-art molecular surface visualization techniques [18] remove critical information such as channels and fill the resultant holes for using spherical parameterization algorithms which can induce considerable angle and area distortion; higher genus surfaces can induce considerably higher area and angle distortion due to higher loss of interior information. We implemented the two force-directed and parameter-based spherical parameterization algorithms, and the two Lambert equal-area and Plate Carrée map projection algorithms, presented in Krone et al. [18]. The results on two molecular genus 0 and genus 1 (with channel removal and hole-filling) surface datasets are shown in Figure 10. The map projections in Figures 10c and 10f were used to create space time cube visualizations [18] to visualize the cavities on the boundaries without the need for animation. LMap can similarly be used to visualize the cavities while preserving the geometric context and without removing the critical channel information, as shown in Figure 11.

5 CONCLUSION AND FUTURE WORK

In this work, we presented LMap, a technique to conformally parameterize and deform selected regions of interest on an arbitrary surface. We achieved this using the robust and theoretically-sound extrinsic Ricci flow. We also guarantee the existence of a local map for a selected ROI on an arbitrary surface via the use of the dynamic Ricci flow. The effectiveness and efficacy of our LMap technique was then demonstrated in three challenging use cases: (1) multimodal brain visualization, (2) optimal coverage of virtual colonoscopy flythroughs, and (3) molecular surface visualization.

In the future, we will incorporate area-preserving techniques into LMap using optimal mass transport for balancing the resultant area- and angle-distortions. Quasi-conformal parameterization [35] allows bounded conformality distortion. In this direction, we will explore extending LMap to quasi-conformal LMaps. Since we can locally parameterize a selected ROI without requiring any global information or spherical constraints [26], we will also validate our approach in the context of decal mapping. Since LMap allows multimodal brain visualization in the native space, we will integrate subcortical structures into these LMap visualizations for a more comprehensive picture of the brain. In the context of VC, we will conduct a formal study with expert radiologists to test LMap in VC. For molecular surface visualization, we will introduce visualizations to study the evolution of the boundary cavities in place with exploded LMap views of these cavities, and

will combine these with context-preserving maps [19] to visualize the interior cavities, simultaneously.

ACKNOWLEDGMENTS

The VC datasets are courtesy of Stony Brook University Hospital (SBUH) and Dr. Richard Choi, Walter Reed Army Medical Center. The brain datasets are courtesy Human Connectome Project, WU-Minn Consortium (PIs: David Van Essen and Kamil Ugurbil) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. The molecular imaging data is courtesy Dr. Michael Krone of Visualization Research Center, University of Stuttgart, Germany. We would like to thank Dr. Matthew Barish and Dr. Kevin Baker of SBUH and Dr. Hoi-Chung Leung of SBU Neuroscience Department for their help in this project. This work has been partially supported by the National Science Foundation grants IIP1069147, CNS1302246, IIS1527200, NRT1633299, CNS1650499, and the Marcus Foundation.

REFERENCES

- [1] N. Aigerman and Y. Lipman. Orbifold Tutte embeddings. *ACM Transactions on Graphics*, 34(6):190:1–190:12, 2015.
- [2] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis. Conformal geometry and brain flattening. *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 271–278, 1999.
- [3] A. V. Bartolí, R. Wegenkittl, A. König, and E. Gröller. Nonlinear virtual colon unfolding. *IEEE Visualization*, pages 411–420, 2001.
- [4] A. V. Bartolí, R. Wegenkittl, A. König, E. Gröller, and E. Sorantin. Virtual colon flattening. *Data Visualization*, pages 127–136, 2001.
- [5] H. M. Berman, J. Westbrook, Z. Feng, G. Gilliland, T. N. Bhat, H. Weissig, I. N. Shindyalov, and P. E. Bourne. The protein data bank. *Nucleic Acids Research*, 28(1):235–242, 2000.
- [6] P. T. Choi, K. C. Lam, and L. M. Lui. Flash: Fast landmark aligned spherical harmonic parameterization for genus-0 closed brain surfaces. *SIAM Journal on Imaging Sciences*, 8(1):67–94, 2015.
- [7] K. Crane, U. Pinkall, and P. Schröder. Spin transformations of discrete surfaces. *ACM Transactions on Graphics*, 30(4):104, 2011.
- [8] S. F. Gibson. Constrained elastic surface nets: Generating smooth surfaces from binary segmented data. *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 888–898, 1998.
- [9] X. Gu, R. Guo, F. Luo, J. Sun, and T. Wu. A discrete uniformization theorem for polyhedral surfaces II. *arXiv preprint arXiv:1401.4594*, 2014.
- [10] X. Gu, Y. Wang, T. F. Chan, P. M. Thompson, and S.-T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transactions on Medical Imaging*, 23(8):949–958, 2004.
- [11] K. C. Gurijala, R. Shi, W. Zeng, X. Gu, and A. Kaufman. Colon flattening using heat diffusion Riemannian metric. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2848–2857, 2013.
- [12] S. Haker, S. Angenent, A. Tannenbaum, and R. Kikinis. Nondistorting flattening maps and the 3-D visualization of colon CT images. *IEEE Transactions on Medical Imaging*, 19(7):665–670, 2000.
- [13] J. Hass and P. Koehl. How round is a protein? Exploring protein structures for globularity using conformal mapping. *Frontiers in Molecular Biosciences*, 1:26, 2014.
- [14] W. Hong, F. Qiu, and A. Kaufman. A pipeline for computer aided polyp detection. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):861–868, 2006.
- [15] M. K. Hurdal and K. Stephenson. Cortical cartography using the discrete conformal approach of circle packings. *Neuroimage*, 23:S119–S128, 2004.
- [16] L. Ju, M. K. Hurdal, J. Stern, K. Rehm, K. Schaper, and D. Rottenberg. Quantitative evaluation of three cortical surface flattening methods. *Neuroimage*, 28(4):869–880, 2005.
- [17] D. G. Kontopoulos, D. Vlachakis, G. Tsiliki, and S. Kossida. Struc-tuprint: a scalable and extensible tool for two-dimensional representation of protein surfaces. *BMC Structural Biology*, 16(1):4, 2016.


- [18] M. Krone, F. Frieß, K. Scharnowski, G. Reina, S. Fademrecht, T. Kulschewski, J. Pleiss, and T. Ertl. Molecular surface maps. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):701–710, 2017.
- [19] J. Marino, W. Zeng, X. Gu, and A. Kaufman. Context preserving maps of tubular structures. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):1997–2004, 2011.
- [20] S. Nadeem and A. Kaufman. Multimodal brain visualization. *SPIE Medical Imaging*, pages 97881Y–97881Y, 2016.
- [21] S. Nadeem, J. Marino, X. Gu, and A. Kaufman. Corresponding supine and prone colon visualization using eigenfunction analysis and fold modeling. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):751–760, 2017.
- [22] S. Nadeem, Z. Su, W. Zeng, A. Kaufman, and X. Gu. Spherical parameterization balancing angle and area distortions. *IEEE Transactions on Visualization and Computer Graphics*, 23(6):1663–1676, 2017.
- [23] D. S. Paik, C. F. Beaulieu, R. B. Jeffrey Jr, C. A. Karadi, and S. Napel. Visualization modes for CT colonography using cylindrical and planar map projections. *Journal of Computer Assisted Tomography*, 24(2):179–188, 2000.
- [24] A. A. Polyansky, A. O. Chugunov, P. E. Volynsky, N. A. Krylov, D. E. Nolde, and R. G. Efremov. Preddimer: A web server for prediction of transmembrane helical dimers. *Bioinformatics*, 30(6):889–890, 2014.
- [25] S. J. Rahi and K. Sharp. Mapping complicated surfaces onto a sphere. *International Journal of Computational Geometry & Applications*, 17(04):305–329, 2007.
- [26] A. Rocha, U. Alim, J. D. Silva, and M. C. Sousa. Decal-maps: Real-time layering of decals on surfaces for multivariate visualization. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):821–830, 2017.
- [27] M. F. Sanner, A. J. Olson, and J.-C. Spehner. Reduced surface: an efficient way to compute molecular surfaces. *Biopolymers*, 38(3):305–320, 1996.
- [28] B. Springborn, P. Schröder, and U. Pinkall. Conformal equivalence of triangle meshes. *ACM Transactions on Graphics*, 27(3):77:1–77:11, 2008.
- [29] D. Tosun and J. L. Prince. A geometry-driven optical flow warping for spatial normalization of cortical surfaces. *IEEE Transactions on Medical Imaging*, 27(12):1739–1753, 2008.
- [30] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, K. Ugurbil, W.-M. H. Consortium, et al. The WU-Minn Human Connectome Project: An overview. *Neuroimage*, 80:62–79, 2013.
- [31] G. Wang, G. McFarland, B. P. Brown, and M. W. Vannier. GI tract unraveling with curved cross sections. *IEEE Transactions on Medical Imaging*, 17(2):318–322, 1998.
- [32] Y. Wang, L. M. Lui, T. F. Chan, and P. M. Thompson. Optimization of brain conformal mapping with landmarks. *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 675–683, 2005.
- [33] Z. Wang, Z. Liang, X. Li, L. Li, B. Li, D. Eremina, and H. Lu. An improved electronic colon cleansing method for detection of colonic polyps by virtual colonoscopy. *IEEE Transactions on Biomedical Engineering*, 53(8):1635–1646, 2006.
- [34] F. Yeh. DSI studio, 2014.
- [35] W. Zeng, J. Marino, K. C. Gurijala, X. Gu, and A. Kaufman. Supine and prone colon registration using quasi-conformal mapping. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1348–1357, 2010.
- [36] X. Zhao, Z. Su, X. D. Gu, A. Kaufman, J. Sun, J. Gao, and F. Luo. Area-preservation mapping using optimal mass transport. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2838–2847, 2013.



Saad Nadeem is a PhD candidate in the Computer Science department, Stony Brook University. He received his BSc Honors in Computer Science and Mathematics from School of Science and Engineering, Lahore University of Management Sciences, Pakistan. His research interests include computer vision, computer graphics, and visualization.



Xianfeng Gu received the Ph.D. degree in computer science from Harvard University, Cambridge, MA, USA, in 2003. He is an associate professor of Computer Science and the Director of the 3D Scanning Laboratory with the Department of Computer Science at Stony Brook University, Stony Brook, NY, USA. His current research interests include computer vision, graphics, geometric modeling, and medical imaging. His major works include global conformal surface parameterization in graphics, tracking and analysis of facial expression in vision, manifold splines in modeling, brain mapping and virtual colonoscopy in medical imaging, and computational conformal geometry. He won the U.S. National Science Foundation CAREER Award in 2004.



Arie E. Kaufman is a Distinguished Professor and Chairman of the Computer Science Department, Director of the Center of Visual Computing (CVC), and Chief Scientist of the Center of Excellence in Wireless and Information Technology (CEWIT) at Stony Brook University. He has conducted research for over 40 years in visualization and graphics and their applications, has published more than 300 refereed papers, has presented more than 20 invited keynote talks, has been awarded/filed more than 40 patents, and has been PI/co-PI on more than 100 grants. He was the founding Editor-in-Chief of IEEE Transaction on Visualization and Computer Graphics (TVCG), 1995-1998. He is a Fellow of IEEE, a Fellow of ACM, the recipient of the IEEE Visualization Career Award (2005), and was inducted into the Long Island Technology Hall of Fame (2013). He received his PhD in Computer Science from Ben-Gurion University, Israel, in 1977.