

CORTICAL SURFACE FLATTENING: A QUASI-CONFORMAL APPROACH USING CIRCLE PACKINGS

MONICA K. HURDAL¹, KEN STEPHENSON², PHILIP L. BOWERS¹,
DE WITT L. SUMNERS¹, AND DAVID A. ROTTENBERG^{3,4}

¹Department of Mathematics, Florida State University,
Tallahassee, FL 32306-4510, U.S.A.

²Department of Mathematics, University of Tennessee,
Knoxville, TN 37996-1300, U.S.A.

³PET Imaging Center, VA Medical Center, Minneapolis, MN 55417, U.S.A.

⁴Departments of Neurology and Radiology, University of Minnesota,
Minneapolis, MN 55455, U.S.A.

Address for correspondence: Dr. Monica K. Hurdal, Department of Mathematics,
Florida State University, Tallahassee, FL 32306-4510, U.S.A.
Fax: (+1 850) 644-4053 Email: mhurdal@math.fsu.edu

Running Title: Quasi-Conformal Flat Maps From Circle Packings

Key Words: circle packing; conformal map; cortical flat map; mapping; Riemann
Mapping Theorem; surface flattening

ABSTRACT

Comparing the location and size of functional brain activity across subjects is difficult due to individual differences in folding patterns and functional foci are often buried within cortical sulci. Cortical flat mapping is a tool which can address these problems by taking advantage of the two-dimensional sheet topology of the cortical surface. Flat mappings of the cortex assist in simplifying complex information and may reveal spatial relationships in functional and anatomical data that were not previously apparent. Metric and areal flattening algorithms have been central to brain flattening efforts to date. However, it is mathematically impossible to flatten a curved surface in 3-space without introducing metric and areal distortion. Nevertheless, the Riemann Mapping Theorem of complex function theory implies that it is theoretically possible to preserve conformal (angular) information under flattening. In this paper we present a novel approach for creating flat maps of the brain that involves a computer realization of the 150-year-old Riemann Mapping Theorem. This approach uses a circle packing algorithm to compute an essentially unique (i.e. up to Möbius transformations), discrete approximation of a conformal mapping from the cortical surface to the plane or the sphere. Conformal maps are very versatile and offer a variety of visual presentations and manipulations. Maps can be displayed in three geometries: the Euclidean and hyperbolic planes, and the sphere. A wide variety of Möbius transformations can be used to zoom and focus the maps in a particular region of interest. Conformal maps are mathematically unique and canonical coordinate systems can also be specified on these maps. Although conformal maps do not attempt to preserve linear or areal information, locally they appear Euclidean. Conformal information allows shape to be preserved. In this paper we describe our approach and present some of advantages of conformal flattening using circle packings. We discuss the notion of a conformal structure on a surface, and describe the three geometries of constant curvature surfaces where our maps reside, as well as classical conformal automorphisms (Möbius transformations) of these surfaces. We describe how circle packing can be used to obtain quasi-conformal mappings of surfaces and demonstrate the advantages of this approach by producing quasi-conformal flat maps with data from the Visible Man and from an MRI volume of the human cerebellum.

1. INTRODUCTION

The cortex of humans and other primates is a highly convoluted surface. The folds (gyri) and fissures (sulci) of the brain vary in size and position from person to person. This variability has made it difficult for medical researchers to analyze and compare functional regions of the brain within and between subjects.

Non-invasive anatomical and functional data are available from a variety of modalities, including magnetic resonance imaging (MRI), functional MRI (fMRI) and positron emission tomography (PET) (see Toga and Mazziotta (1996) for a compilation of various methods). A common approach to visualizing functional data is to project a focus of activation on a two-dimensional cross-section or slice of a brain volume. However, this approach suffers from a number of disadvantages due to the highly folded structure of the cerebral cortex. Activated foci that appear close together on a two-dimensional brain slice or 3D surface rendering may be quite far apart when visualized on the unfolded cortical surface. In a given individual, foci are often buried within cortical sulci and appear in a number of discrete slices, making it difficult to compare multiple foci simultaneously. Comparing the location and size of functional activity across subjects is difficult due to individual differences in folding patterns (Rehm *et al.*, 1998).

To address some of these problems, a number of methods have been implemented that take advantage of the two-dimensional sheet topology of the cortical surface. These methods begin by reconstructing and then flattening the cortical surface. Flat mappings of the cortex assist in simplifying complex information and may reveal spatial relationships in functional and anatomical data that were not previously apparent. Traditionally, the complex surface of the cerebral cortex was reconstructed by tracing contours from histological sections. Wire frames or sheets were created from the contours and these were then aligned and stacked to reconstruct the surface (Rosa *et al.*, 1997). More recently, Van Essen and colleagues (Drury *et al.*, 1996; Drury and Van Essen, 1997; Van Essen and Drury, 1997; Van Essen *et al.*, 1998) and Dale and colleagues (Dale and Sereno, 1993; Dale *et al.*, 1999; Fischl *et al.*, 1999) have developed computational tools for this purpose. These tools generally work by attempting to reduce the metric or areal distortion between the original surface and the flattened surface.

Computational cortical surface reconstruction from MRI scans yields a polygonal mesh representing the surface of the grey matter. This mesh generally contains topological errors, such as holes and handles, which must be corrected before surface flattening can begin. In order to reduce distortions during flattening, cuts are introduced in the corrected mesh. The surface is then flattened by applying functions that adjust the edge lengths forming the polygonal mesh, allowing the surface to unfold. Drury *et al.* (1996) use longitudinal and torsional forces so that linear and angular distortions are reduced while unfolding the surface. Fischl *et al.* (1999) use the gradient of a function incorporating geodesic distance and area so that linear and areal distortions are reduced. These types of metric and areal flattening algorithms have been central to brain flattening efforts to date.

Flat mapping traces at least to Ptolemy's efforts to represent the earth's surface on a flat sheet of paper. Shortcomings were evident from antiquity, and attempts to preserve

one or another geometric property underlay the myriad types of maps developed down the ages. The mathematical issues are now much clearer and one can prove that it is *impossible* to flatten curved 3D surfaces without introducing metric and areal distortions (Polya, 1968). On the other hand, the Riemann Mapping Theorem of 1852 from complex function theory showed that it *is possible* to preserve angular — technically, conformal — information (Ahlfors, 1966). Knowing such maps exist is not enough as coarse approximations were even impossible to calculate until recently, particularly for surfaces as highly curved and complicated as the brain.

In this paper we present a novel approach for creating flat maps of the brain that involves a computer realization of the 150-year-old Riemann Mapping Theorem. This approach uses a circle packing algorithm to compute an essentially unique discrete approximation of a conformal map that carries the cortical surface into one of the classical geometric spaces, the sphere, the plane, or the hyperbolic plane. Well-understood self-maps of the target spaces, classical functions called Möbius transformations, preserve conformal integrity yet allow one to navigate freely in the flat maps. The canonical nature of the maps and the versatility provided by transformations are the principal advantages of conformal flattening, and these persist in our discrete approximations. For example, our hyperbolic flat maps have the potential for establishing self-consistent coordinate systems since they all reside in a common setting, a disc, yet with Möbius transformations one can instantly bring any region of interest into focus.

We illustrate our flat-mapping methods with two cortical surfaces: a human cerebellum and a human cerebrum. We begin in Section 2 with a brief description of the cortical surface isolation techniques that we employed to obtain the triangulated cortical surface reconstructions we use. In Section 3 we describe the notion of a conformal structure on a surface, define conformal maps, and discuss the three classical surfaces of constant curvature where our maps reside and their conformal automorphisms (Möbius transformations). In Section 4 we describe how circle packing can be used to obtain quasi-conformal mappings of surfaces, and in Section 5 we demonstrate the advantages of this approach by producing quasi-conformal flat maps with data from the Visible Man (The National Library of Medicine, 2000) and from an MRI volume of the human cerebellum. Here we illustrate the versatility provided by the Möbius transformations. We conclude by demonstrating the versatility of our approach and provide a discussion of what we see as some of advantages of conformal flattening using circle packings.

2. CORTICAL SURFACE ISOLATION

In order to create a flattened map of a surface, we require a discrete representation of that surface. For cortical data, the process of obtaining such a surface is an involved and often tedious procedure. Each step in the process is the subject of numerous publications which will not be discussed in detail here.

We use triangulated meshes as discrete representations of our surfaces. All current flattening approaches, including our own, require a topologically correct surface, i.e. a piecewise flat linear surface that triangulates a 2-dimensional sphere or disc. The flat triangular faces are connected along edges. Each edge of the mesh is an interior edge (contained in exactly two triangles) or a boundary edge (contained in exactly one

triangle). If there is a boundary, it has only one boundary component; that is, there is a single closed chain of boundary edges forming the boundary. Formally, this mesh is an oriented surface that is topologically equivalent to a disc (when a boundary exists) or a sphere (when there is no boundary). We illustrate our flat-mapping method using two cortical surfaces.

Human cerebellar surface. A high-resolution T1-weighted MRI volume was acquired (Holmes *et al.*, 1996) and a strategy for isolating the cerebellum in a consistent manner across subjects was developed (Rehm *et al.*, 2000). The cerebrum, brainstem and cerebellar peduncles were stripped away, producing a cerebellar volume defined by a plane parallel to the posterior commissure-obex line and orthogonal to a plane passing through the vermal midline. The result is an isolated cerebellum with angled cuts of the brainstem peduncles analogous to those produced by dissection (Divernoy, 1995). This volume was subsequently parcellated according to Schmahmann *et al.* (1999) and labels (twelve lobes plus white matter) were assigned to each voxel.

To facilitate the production of a simply-connected smooth surface, the cerebellar volume was heavily smoothed and an isosurface triangulated mesh of the surface was created using a marching cubes algorithm (Lorensen and Cline, 1987; Schroeder *et al.*, 1998). The marching cubes algorithm is known to produce topological defects. These were corrected semi-automatically using in-house software. The resulting surface, a topological 2-sphere composed of 56,676 triangles and 28,340 vertices, was parcellated as described above (Figure 1).

=====
Figure 1 about here.
=====

To create a planar flat map of a topological 2-sphere, a boundary must be introduced into the surface to act as the boundary of the 2-disc under flattening. A boundary corresponding to the white matter cut-plane and filled-in fourth ventricle was introduced where the brainstem attaches to the cerebellum. Introducing the boundary resulted in a surface topologically equivalent to a 2-disc. This surface and the uncut surface are used in our surface flattening procedure.

Human cerebral surface. This surface corresponds to the right hemisphere of the Visible Man data set (The National Library of Medicine, 2000). A surface reconstruction of this data was created by David Van Essen and Heather Drury at Washington University and is described in Drury *et al.* (1996) and Drury and Van Essen (1997). This surface, illustrated in Figure 2, was selected to demonstrate that our flattening strategy applies equally well to any cortical surface. It contains 103,845 triangles and 52,360 vertices. The surface was parcellated into anatomically defined lobes and four cuts have been introduced along various sulci to create a surface boundary. These extra cuts were introduced by Van Essen and Drury in an effort to reduce length and areal distortion in their flattening procedure.

=====
Figure 2 about here.
=====

3. CONFORMAL MAPS AND SURFACES OF CONSTANT CURVATURE

In this section we discuss surface geometry, conformal mapping, and our target surfaces, the sphere, plane, and disc. This is an overview: the interested reader will find additional details in the Appendix.

In describing a surface S in 3-space, preliminary considerations include topology, orientation, and smoothness. The surfaces of interest here are topological spheres or, if they have boundary, topological discs (i.e. connected, oriented, genus zero and with, respectively, no boundary or one boundary component) and can be thought of as a wrinkled rubber sphere or a wrinkled rubber sheet, respectively. Each cortical surface has an obvious “outward” direction giving it an orientation. Regarding smoothness, our cortical “surfaces” are merely finite collections of data points from which we create triangulated polyhedra. Treating these as *piecewise smooth* surfaces gives us a well-developed language and full range of fundamental mathematical concepts and theory; in particular, our surfaces inherit from \mathbb{R}^3 a *Riemannian metric*.

The Riemannian metric on S determines the three main structures that are pertinent to flattening: 1) metric structure, meaning lengths and areas; 2) curvature structure, related to peaks, valleys, and folds of S ; and 3) conformal structure, reflecting angles between curves in S . The mathematical notions of *length*, *area*, *curvature*, and *angle* on S all fit naturally with our native intuitions about these quantities. Note that we use what is known as *Gaussian curvature* because it is intrinsic to the surface S and doesn’t depend on knowing how S lives in 3-space.

The goal in flattening is to move data from a cortical surface to a potentially more useful setting. A *mapping* between surfaces S_1 and S_2 is a one-to-one function $f : S_1 \rightarrow S_2$ identifying each point p of S_1 with a corresponding point $q = f(p)$ in S_2 ; the function f (alternately, the image $f(S_1)$ in S_2) is called a *map* of S_1 . The target surfaces for our maps are the three classical geometric surfaces represented by the (Riemann) sphere $\mathbb{S}^2 = \{(x, y, z) \in \mathbb{R}^3 : x^2 + y^2 + z^2 = 1\}$, the *Euclidean* plane \mathbb{R}^2 , and the *hyperbolic* plane modeled as the unit disc $\mathbb{D} = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 < 1\}$. These spaces have well-known Riemannian metrics of constant Gaussian curvature $+1$, 0 , and -1 , respectively. We generalize the term *flat map* of S to mean a map from S into any one of these spaces of constant curvature.

Value is often attached to the metric structure S inherits from 3-space, and a mapping which preserves distance is said to be an *isometry*. However, a surface with non-constant Gaussian curvature cannot be mapped isometrically to a surface of constant Gaussian curvature; that is, *every flat mapping of a cortical surface necessarily introduces metric distortion*.

On the other hand, each of our surfaces S also inherits a complementary but less familiar *conformal structure* related to angular data. A mapping is said to be *conformal* if it preserves the angles (magnitude and orientation) between any two intersecting curves in S . The Riemann Mapping Theorem (RMT) of 1852 (Ahlfors, 1966) asserts that *for every simply connected surface S having a conformal structure there is a conformal map from S onto precisely one of \mathbb{S}^2 , \mathbb{R}^2 , or \mathbb{D}* . In our setting, every cortical surface can be mapped conformally onto \mathbb{S}^2 or \mathbb{D} depending on whether it is a topological

sphere or disc, respectively. The RMT also asserts that such mappings are *essentially* unique, meaning up to Möbius transformations, a well-defined class of normalizations discussed in the Appendix. As for Euclidean maps, if S is a topological disc then it has an astounding array of conformal maps; for example, one can deduce from the RMT that S can be mapped conformally onto *any* region in \mathbb{R}^2 that is bounded by a simple closed curve.

So the first advantage of conformal flat mapping is the guarantee of existence and conditions for uniqueness. Additional advantages lie with the geometries of the classical surfaces in which the maps reside. These surfaces are conveniently nested: $\mathbb{D} \subset \mathbb{R}^2$, and \mathbb{R}^2 is routinely identified, under stereographic projection, with \mathbb{S}^2 punctured at the south pole (see the Appendix). Moreover, it is convenient to identify \mathbb{R}^2 with the *complex* plane \mathbb{C} , so every point $(x, y) \in \mathbb{R}^2$ corresponds with a complex number $z = x + iy$. This gives us the advantage of a complex arithmetic which can be used in any of the three settings. Each target surface \mathcal{T} has a rich family of (*conformal*) *automorphisms*, one-to-one conformal maps from \mathcal{T} onto itself; these form a mathematical *group* under composition denoted by $\text{Aut}(\mathcal{T})$.

Looking at the surfaces separately we can highlight their main geometric features; further details are given in Table 1 and in the Appendix.

The Sphere. \mathbb{S}^2 has constant Gaussian curvature $+1$. Circles are intersections of planes with \mathbb{S}^2 ; geodesics are arcs of great circles. The automorphism group, $\text{Aut}(\mathbb{S}^2)$, is precisely the group \mathcal{M} of all Möbius transformations of \mathbb{S}^2 (see Table 1). Though automorphisms map circles to circles, they do not respect circle centers and geodesics. Any triple of points of \mathbb{S}^2 can be mapped to any other triple by a unique element of $\text{Aut}(\mathbb{S}^2)$. In Figure 3(a), for instance, the north and south poles, N and S, and an equatorial point E are marked and a familiar latitude/longitude grid is imposed. Automorphisms yield Figures 3(b) and 3(c). In 3(b), S and N remained fixed while E was shifted towards N, while in 3(c), S, N, and E were all moved to new locations.

=====
 Figure 3 about here.
 =====

The Plane. The Euclidean plane, identified now as \mathbb{C} , has constant Gaussian curvature zero. $\text{Aut}(\mathbb{C})$ consists precisely of the *affine* maps (rotations, dilations, and/or translations), so each automorphism will map circles to circles, centers to centers, and geodesics to geodesics. The complex plane can be mapped to the surface of the sphere using stereographic projection. This maps both circles and straight lines of \mathbb{C} to circles of \mathbb{S}^2 .

The Disc. The *hyperbolic plane* is a geometric surface of constant Gaussian curvature -1 . There are many ways to represent it concretely; for us the most convenient is the Poincarè disc model. Its point set is the open unit disc in \mathbb{C} , namely $\mathbb{D} = \{z : |z| < 1\}$, and distances between two points are measured in the metric ρ defined by $\rho(z, w) = \frac{1}{2} \log \left(\frac{|1 - z\bar{w}| + |z - w|}{|1 - z\bar{w}| - |z - w|} \right)$. The points of the unit circle, $\partial\mathbb{D}$, while *not in* the hyperbolic plane, can be regarded as an “ideal” boundary, and hyperbolic distances grow as one approaches that boundary. Thus points z, w which appear to our

“Euclidean” eyes as being close to one another can, if near $\partial\mathbb{D}$, be separated by a huge hyperbolic distance. In fact, any path running from a point of \mathbb{D} to an ideal boundary point will have infinite hyperbolic length (see Figure 4).

=====
 Figure 4 about here.
 =====

Hyperbolic circles correspond with Euclidean circles lying in \mathbb{D} (though hyperbolic centers and radii are distinct from Euclidean centers and radii). Hyperbolic geodesics correspond with arcs of Euclidean circles which meet the unit circle $\partial\mathbb{D}$ in right angles. Any Euclidean circle which is internally tangent to $\partial\mathbb{D}$, called a *horocycle*, is treated as a hyperbolic circle of infinite radius with the point of tangency as its (ideal) center. Unlike in the other geometries, all automorphisms of \mathbb{D} are *isometries*, so they preserve hyperbolic circles, circle centers, and geodesics.

Figure 5 illustrates this geometry. Figure 5(a) displays several objects: a shaded disc of hyperbolic radius 0.4 centered at the origin, a triple of mutually tangent circles (including a horocycle) and the (hyperbolic) triangle formed by their centers, and a centered polar-coordinate-style reference grid. Figures 5(b) and 5(c) show these same objects after applying the automorphism $\phi : z \mapsto (3z + 1)/(3 + z)$ once and twice, respectively. In each image the shaded disc has the same hyperbolic radius, though our Euclidean eyes see it as getting progressively smaller as it approaches the boundary. Also the grid remains an orthogonal grid, but with a new “pole”.

=====
 Figure 5 about here.
 =====

Summary. The three classical geometries form a unified and nested hierarchy. They are nested as sets, $\mathbb{D} \subset \mathbb{C} \subset \mathbb{S}^2$; their automorphisms are nested as subgroups of \mathcal{M} , $\text{Aut}(\mathbb{D}) < \text{Aut}(\mathbb{C}) < \text{Aut}(\mathbb{S}^2) = \mathcal{M}$; and both the automorphisms and the inclusion maps preserve circles. In particular, any set C which is a circle in one space will necessarily be a circle in any of the other spaces which happens to contain it. Also, all the geometries are “locally Euclidean”, meaning that at high magnification, the neighborhood of any point looks Euclidean — small circles look like Euclidean circles, geodesics look like Euclidean straight lines, and so forth. In other words, these geometries all look locally like our familiar Euclidean world. Various properties of these geometric surfaces are summarized in Table 1; see the Appendix for details.

=====
 Table 1 about here.
 =====

4. COMPUTING APPROXIMATIONS TO CONFORMAL MAPS VIA CIRCLE PACKING

Conformal maps are guaranteed to exist by the RMT, but are impossible to compute precisely — and until recently, have in fact not even been susceptible of approximation. However, an area of mathematics known as *circle packing*, introduced in 1985 (Thurston, 1985), now provides a computational and theoretical framework we can exploit. The

mathematical focus in this paper is on practical *numerical* approximations of conformal maps *via* circle packing. The resulting “discrete conformal” maps have a controlled but non-negligible level κ of conformal distortion — technically they are κ -*quasiconformal*, as we describe in the Appendix. We note that while reducing this distortion is desirable, the *discrete* conformal geometry of circle packing has such strong parallels to *classical* conformal geometry in both theory and intuition that to a large degree the advantages of conformal maps are already present in our discrete versions.

A *circle packing* is a configuration of circles with a specified pattern of tangencies. Figure 6 shows some examples, but for us the underlying pattern will be associated with a cortical surface S . In particular, S is presented as a triangular mesh that is topologically equivalent to a sphere or disc. It is important to isolate the *pattern* of this mesh from its *geometry*, so we describe the triangulation in terms of its combinatorics K and its geometric realization V , where K is a (simplicial) complex representing the connectivity of the vertices, edges, and faces while V is a set of vertex positions (points in \mathbb{R}^3) defining the shape of the mesh. We write $S = (K, V)$.

=====
 Figure 6 about here.
 =====

Given K , a circle packing P for K in one of our target surfaces \mathcal{T} is a collection $\{c_v\}$ of circles in \mathcal{T} , one for each vertex v of K , so that c_v is tangent to c_u whenever $\langle v, u \rangle$ is an edge of K and so that a triple $\langle c_v, c_u, c_w \rangle$ of mutually tangent circles is positively oriented in \mathcal{T} whenever $\langle v, u, w \rangle$ is a positively oriented face of K . Such a circle packing gives us a new triangulated surface S' lying in \mathcal{T} where $S' = (K, V')$, i.e. the same combinatorics but a new set of vertex positions determined by the centers of the circles. The induced triangulation S' of a packing is shown in Figure 7. The new surface S' is the *flat map* of the original surface S . Since S and all the associated packings share the same combinatorics, as encoded in K , each vertex of S corresponds to a vertex v of K and hence to a circle c_v of P . Thus the three vertices defining any face of S may be identified with the three vertices defining a triangle in S' , yielding a surface map $f : S \rightarrow S'$. A map defined in this way will be called a *discrete conformal map*.

=====
 Figure 7 about here.
 =====

Of course, our discrete conformal maps depend on the existence of circle packings for given (extremely complicated) patterns K . Existence follows from results of Koebe, Andreev, and Thurston (Koebe, 1936; Andre’ev, 1970; Thurston, 1997). If one treats K as a type of *discrete* conformal structure on S , these may be formulated in parallel with the classical theory. Thus the Discrete Riemann Mapping Theorem (DRMT) asserts that *for every simply connected triangulated surface S there is a discrete conformal map from S onto precisely one of \mathbb{S}^2 or \mathbb{D}* . More precisely, if K is a topological sphere then there exists a circle packing P_K for K lying in \mathbb{S}^2 and P_K is unique up to automorphisms of \mathbb{S}^2 , while if K is a topological disc, then there exists a circle packing P_K for K

lying in the hyperbolic plane \mathbb{D} , the circles associated with boundary vertices of K are horocycles, and P_K is unique up to automorphisms of \mathbb{D} .

The packing P_K is called the *maximal* packing for K , and as noted is essentially unique. When K is a topological disc, however, prescribing boundary radii leads to a plethora of other discrete conformal maps. Indeed, if v_1, \dots, v_n denote the boundary vertices of K , then the Circle Packing Theorem (CPT) (Beardon and Stephenson, 1990) asserts that given any assignment of positive numbers r_1, \dots, r_n there exists a unique (up to Euclidean isometry) circle packing P in \mathbb{R}^2 such that for each $i, i = 1, \dots, n$, the boundary circle c_{v_i} of P has the assigned radius r_i . The flexibility provided by these Euclidean packings will be demonstrated in examples later.

The practical question of computing the circle packings guaranteed by the DRMT occupies the remainder of this section, with further details provided in the Appendix. The key is the collection R of radii for the circles, termed a *packing label* for K . Once R is known, it is straightforward to construct (i.e. lay out) P .

The *packing process* refers to the methods used to compute (i.e. to approximate) R , and the key is the extensive system of “flat” local compatibility conditions R must satisfy. The curvature of a piecewise flat surface S is concentrated at its vertices. Any interior vertex v of S has a chain of contiguous neighboring vertices that form the triangular faces surrounding v . To “flatten” the surface at v thus requires adjustment of these triangles so that the angles they form at v sum to 2π (360 degrees). If the adjusted triangles are determined by circles and if their radii are recorded in a label $R = \{r_v\}$, then trigonometry provides a *packing condition* for flatness at v . In the Euclidean setting, this condition is

$$(1) \quad \sum_{\langle v,u,w \rangle} \arccos \left\{ \frac{(r_v + r_u)^2 + (r_v + r_w)^2 - (r_u + r_w)^2}{2(r_v + r_u)(r_v + r_w)} \right\} = 2\pi,$$

where this sum is over all faces $\langle v, u, w \rangle$ containing v . Analogous formulae apply in the other geometries (see Table 1). R is a packing label if and only if the packing condition holds at every interior v . There is (typically) no packing condition for boundary vertices, which accounts for the extra degrees of freedom a boundary provides.

Computation of R uses an iterative process proposed by Thurston (1985) and refined by recent work of Collins and Stephenson (2001). The idea behind the algorithm is simple and elegant. A failure of equality in Equation 1 for a particular v can be remedied by decreasing r_v , the radius of v , if the angle sum is too small (i.e. less than 2π) or by increasing r_v if the angle sum is too large (i.e. greater than 2π). Figures 8 and 9 illustrate how decreasing and increasing, respectfully, the radius of a circle can lead to the correct packing condition with its neighbors.

=====
 Figure 8 about here.
 =====

=====
 Figure 9 about here.
 =====

In practice, one starts a packing computation by assigning desired radii to boundary vertices (if any) and arbitrary labels (i.e. putative radii) to interior vertices. One then repeatedly readjusts the labels of the interior vertices, one at a time as described above, until the packing conditions of (1) are (approximately) satisfied for all interior v . The boundary radii (if any) do not change. There is a comprehensive theory guaranteeing that the iterative scheme converges to a unique packing label R . Using the values in R as radii, one lays out the circles to get the associated circle packing P in the target surface \mathcal{T} . The packing in turn determines the associated conformal map from S to \mathcal{T} , giving a quasi-conformal flat map of S .

A few final comments are in order. There is no packing algorithm known to work in the spherical setting: all spherical packings are computed in \mathbb{C} or \mathbb{D} and then stereographically projected to \mathbb{S}^2 , as described in Section 5. The iterative computations we have described are purely *numeric*, not geometric, meaning there is no coherent circle packing possible until a sufficiently accurate approximation of the final packing label is reached. Also note that although the packing algorithm involves “local” adjustments, each depending only on a vertex v and its immediate neighbors, the theory tells us that the consequences are ultimately global. This is, in fact, the discrete expression of the rigidity inherent in all conformal structures and underlies the utility of the discrete methods. Our group has adapted a software package called `CirclePack` created by Stephenson that computes the packing radii in any of the three geometries and provides an interface for displaying and manipulating the resulting circle packings in any of the three geometries.

5. CORTICAL FLAT MAPS

We illustrate the creation and manipulation of flat maps using the concrete surface triangulations introduced in Section 2. Recall that after preprocessing, each complex K is a topologically correct sphere or disc. One must exploit the connection with the original 3D triangulated surface S , each point in the flat map being associated with a point of S and *vice versa*. Fundamental information resides with the surface — landmarks, standard region demarkations, surface distances, sulci, gyri, surface curvature, etc. Certain aspects of this information can be transferred to the flat map, typically using common color coding. A region or curve can be marked on the map, but relevant data such as surface areas and lengths, curvatures, and other metric information, must always be computed based on the connection to S .

The major computational effort in flattening with circle packings involves approximation of the packing label for the underlying graph K . For small surfaces (under 50,000 vertices), this computation takes less than an hour on a Pentium II 400 MHz PC. As the surface size increases, the computation times increases substantially. This computation is a one-time task for a given data set and subsequent manipulations and transformations occur in real-time. The visual presentation of a circle packing flat map uses either the circles of the packing or the faces of the geometric triangulation induced by the packing (see Figure 7). Color is used to encode additional information, such as anatomical or functional data. The original 3D surface and flat maps in any of the

geometries can be viewed simultaneously in neighboring windows. Points or regions chosen on one can be highlighted on the other.

Euclidean Maps. The surface must be a topological disc to be flattened in \mathbb{R}^2 . When using this scheme for flattening, we assign to each boundary vertex w a label equal to the average of half the lengths of the two boundary edges containing w on the original surface. Because of the boundary conditions, the resulting flat map will have the Euclidean lengths on the boundary preserved. `CirclePack` then computes the unique packing label R for K having these prescribed labels for the boundary radii. Circle packings roughly preserve conformal, not metric structure. We observe that minimizing metric distortion on the boundary edges does not imply that metric distortion is small in the interior also. In \mathbb{R}^2 transformations of the packings involve the Möbius transformations of $\text{Aut}(\mathbb{C})$, translations, dilations, and rotations, giving the user options for arbitrary real-time renormalizations in `CirclePack`.

We consider the surfaces described in Section 2. The cerebellar surface acquired a boundary when the brainstem was removed: 222 vertices/edges. The visible man data contains a boundary of 873 vertices/edges. A normalization is required before laying out the circles. For the cerebellum we have chosen to place the center of the horizontal fissure at the origin and the base of the primary fissure vertically above the origin. The flat maps for these surfaces are shown in Figure 10.

=====
 Figure 10 about here.
 =====

Hyperbolic Maps. This setting also requires that the surface be a topological disc. We use `CirclePack` to compute the so-called maximal packing P_K by simply assigning the hyperbolic label ∞ to all boundary radii and repacking the interior. The boundary circles become horocycles in the final configuration. The cerebellum we packed in the Euclidean setting is repacked in \mathbb{D} in Figure 11, left; the same normalizations were used for the origin and a point directly above the origin (on the positive y -axis). The cerebrum is repacked and shown in Figure 12, left. Note that the enclosing outer circle represents the boundary of \mathbb{D} (the unit circle) and is not a circle of the packing.

=====
 Figure 11 about here.
 =====

=====
 Figure 12 about here.
 =====

There are two main features to highlight regarding hyperbolic flat maps. First, the final packings all lie in a common setting (i.e. they are all disc-shaped), regardless of data set sizes, normalizations, boundary, any *ad hoc* surface cuts, and so forth. This simplifies one of the primary registration difficulties encountered with flat mappings and presents the mapping in a standard setting consistent with gathering meaningful statistics. Second is the rich group of rigid hyperbolic motions or Möbius transformations, $\text{Aut}(\mathbb{D})$. The map center is of visual importance in \mathbb{D} because there the map

appears Euclidean with little hyperbolic distortion. `CirclePack` provides real-time interaction to bring any interior circle to the origin, allowing the map focus to be changed. Figures 11 and 12 show the effect. A reference coordinate system has been imposed centered on a point p (Figure 11, left). A user interested in the region around the point q simply selects that point with a mouse click, immediately transforming the image to Figure 11, right. The transformed coordinate grid shows the effects, with the grid lines still intersecting each other orthogonally. In other words, the view is akin to that in a light microscope: the area of interest is brought to the center and what appears to our Euclidean eyes as distortion is pushed to the periphery. Recall, however, that the automorphisms are actually isometries, rigid motions in the hyperbolic sense, so these changes in focal point have absolutely no effect on any intrinsic hyperbolic structures needed for computations and statistics.

Spherical Maps. A topological sphere S cannot be mapped into \mathbb{R} or \mathbb{D} without introducing cuts; however, there is an essentially unique conformal map to the sphere \mathbb{S}^2 which we approximate with discrete conformal maps. Since there is no packing algorithm intrinsic to spherical geometry, we employ a mathematical trick. An arbitrary vertex v_0 and all edges containing it are removed (or punctured) from K , leaving a topological disc K^* . Its maximal circle packing P_{K^*} in \mathbb{D} is computed and then projected stereographically to the sphere; all its boundary circles are tangent to the equator (see Figure 15(f)). The equator is introduced as the circle for v_0 , resulting in a spherical circle packing for the original complex K . This circle packing is then normalized by an appropriate automorphism of \mathbb{S}^2 , after which the choice of v_0 is immaterial.

Consider the data from Figure 1 which is a topological sphere. The automorphism group of \mathbb{S}^2 allows one to choose three points for normalization, typically points for the poles N and S and a point to be placed at E on the equator. Figure 13 displays the spherical packing for the cerebellum with the precentral fissure mapped to N , the center of the horizontal fissure mapped to S and the base of the primary fissure mapped to E . Figure 13(a) is a typical view and Figure 13(b) rotates the sphere to display the location of the brain stem which was the boundary used for the other flat maps. Figure 13(c) illustrates the application of an automorphism.

=====
 Figure 13 about here.
 =====

Comparisons and Coordinate Systems. The normalizations required for displaying circle packings provide a means for imposing canonical coordinate systems on these flat maps. In the case of the Euclidean and hyperbolic maps, two points, such as anatomical landmarks are required, while for spherical maps it is three points (two poles and an equator). The normalizations discussed in the previous sections form the basis for the coordinate grids on the flat maps in Figures 11 and 14.

It is instructive to compare the three flat maps associated with our cerebellar data. In Figure 14 we focus the images from Figures 10, 11 and 13, respectively, on the region near a common point, the base of the horizontal fissure. The coordinate reference grids are also displayed. The boundary where the brain stem was removed has drastically

different shapes in the three settings, but despite that one can see that the local structures in the interior are nearly identical in the three maps. The similarities would be even more striking if we focused in on a smaller region. This interior integrity despite the mapping modality is one of the key consequences of the conformal nature of the maps.

=====
 Figure 14 about here.
 =====

6. VERSATILITY OF THE CIRCLE PACKING APPROACH

Circle packings provide enormous flexibility and versatility for computing flat maps. To demonstrate we will let S be the small, triangulated surface in 3-space pictured in Figure 15(a), which is a simply connected patch from Figure 2. The complex K (211 triangles, 122 vertices including 31 boundary vertices/edges) is a topological disc, so it has a variety of packings, as illustrated by Figure 15. Figure 15(b) displays the maximal packing P_K in \mathbb{D} guaranteed by DRMT. Figures 15(c) and 15(d) are two Euclidean packings of K , each with its induced triangulation for reference. In 15(c) the radii of the boundary circles were set in advance so that the lengths of boundary edges would match their lengths in S . Figure 15(d) illustrates an alternate type of boundary condition; here the carrier was required to form a rectangle with four preselected boundary circles as its corners. Projecting Figures 15(c) and 15(b) to the sphere gives the circle packings of Figures 15(e) and 15(f), respectively; Figure 15(e) shows just the induced triangulation (now spherical) and 15(f) suggests the geometry of stereographic projection.

=====
 Figure 15 about here.
 =====

The circle packing approach admits various novel boundary conditions for flat maps, whether one is mapping the full surface or some chosen fragments. Packing a rectangular shape with designated corner vertices as in Figure 15(d) is one example, and in this case the ratio of the length to the width of the resulting rectangle (its *conformal modulus*) can be shown to be uniquely determined by the conformal structure of S . We do not at this time propose a scientific role for such conformal invariants nor for these alternative packings, but the conformal structure on a surface involves some very subtle “shape” information which may find its uses in the future.

7. DISCUSSION

Our novel approach to flat-mapping cortical surfaces relies on the mathematical theory of conformal maps. This offers a number of advantages over other strategies; namely, guaranteed existence, uniqueness, and a rich theoretical framework. We use *circle packing* methods to define *discrete conformal maps* which both approximate conformal maps and enjoy a parallel discrete theory. In particular, they exist and are unique by the DRMT and the CPT, are computable, and can be manipulated in practice in the same ways that conformal maps can be manipulated in theory. In addition, discrete conformal maps are κ -quasiconformal with bounds on the conformal distortion κ (see the

Appendix). We are also able to refine the circle packing approach to produce closer approximation to the true conformal surface. We are developing new theoretical computational approaches for producing even closer approximations by preserving inversive distance (a hyperbolic invariant) rather than circle tangency.

Conformal maps are versatile and offer a variety of visual presentations and manipulations. Maps can be displayed in three geometries: the Euclidean and hyperbolic planes, and the sphere. Each setting has different advantages depending on the requirements. The three geometries are related and similarities between maps in one geometry and another are readily apparent. In addition, a wide variety of Möbius transformations can be used to zoom and focus the maps in a particular region of interest.

The Euclidean plane has the advantage of familiarity. As with flat maps produced by other researchers (Drury *et al.*, 1996; Fischl *et al.*, 1999), the shape of this map is largely determined by the length and number of edges in the chosen boundary and will vary from map to map. The hyperbolic map, on the other hand, is always the same disc shape, making it easier to compare maps. Although the hyperbolic metric at first seems distorting, Möbius transformations allow the map focus to be changed interactively, relegating the distortion to the map boundary. As with other flat mapping approaches, an initial boundary cut is required to create the Euclidean and hyperbolic maps, but the extraneous cuts along fissures used by other flat mapping methods are **not** required here. Of course, the primary advantage of a spherical map is that it requires no cuts at all.

Canonical coordinate systems can also be specified on these maps by identifying two landmarks for Euclidean and hyperbolic maps, or three landmarks for the sphere. In the case of the hyperbolic and spherical maps, identifying the same landmarks on other hyperbolic or spherical maps automatically puts the maps in register. Additional warping is likely needed if fissures or lobes are to be aligned, but the underlying canonical coordinate system can be carried along.

Conformal maps do not preserve the linear and areal information that may be important in gauging neuronal density and functional activation. But it must be pointed out that even strategies keyed to the metric data can only minimize distortion locally — global measurements will always be done by reference to the original surface. As to any potential significance to conformal information in-and-of-itself, this is still open. Nevertheless, a number of studies have suggested a conformal mapping between the visual field and the visual cortex (Fischer, 1973; Tusa *et al.*, 1978; Schwartz, 1980; Murray, 1989; Schwartz, 1994). Conformal maps appear Euclidean at the local level and they do preserve subtle shape information which could well play a role in cortical studies. Taken along with their practical advantages, conformal maps thus may allow us to better localize functional regions of activation in normal subjects and patients with other hereditary diseases.

8. SOFTWARE AVAILABILITY

The software, `CirclePack`, which was used to create these flat maps is available from <http://www.pet.med.va:8080/incweb>. A topologically correct triangulated surface that is equivalent to a sphere or disc is required before using this software for conformal

flattening. The software currently has been compiled to run under Linux, SunOS and Solaris.

9. ACKNOWLEDGEMENTS

This work and adaption of the `CirclePack` software for neuroscience is supported in part by NIH Human Brain Project grant MH57180. Development of the `CirclePack` software has been supported in part by the National Science Foundation and the Tennessee Science Alliance. The authors would like to acknowledge Kelly Rehm, Kirt Schaper and Josh Stern, VA Medical Center and Departments of Neurology and Radiology, University of Minnesota, Minneapolis. David Van Essen, Heather Drury and James Dickson of Washington University School of Medicine, St. Louis are also gratefully acknowledged for providing the Visible Man data.

APPENDIX A.

We provide background and additional details on the mathematics of surfaces, the classical geometric spaces, conformal maps, and circle packings.

Primer on Surfaces. Our surfaces S are intended to be polyhedral topological spheres or discs defined by triangular meshes. Typical topological problems — edges occurring more than twice, disconnected pieces, etc. — are detected and repaired by examining the surface complex K . So-called *handles* are more difficult, but can be detected using topological invariants; a surface's *Euler characteristic* $\chi(S)$ is defined by $\chi(S) = v - e + t$, where v, e, t are the numbers of vertices, edges, and triangles, respectively, in K (Armstrong, 1983; Massey, 1967). The number $m(S)$ of boundary components can be computed from K . The *genus* $g(S)$, the number of handles, satisfies $\chi(S) = 2 - 2g(S) - m(S)$. Thus, assuming S has at most one boundary component and is topologically correct, S is a topological sphere if and only if $\chi(S) = 2$ and a topological disc if and only if $\chi = 1$. In addition, $g(S) = 0$.

Since our surfaces lie in 3-space and are piecewise flat, they inherit a Riemannian metric from \mathbb{R}^3 . This endows S with the three main structures pertinent to flattening: metric structure, curvature structure and conformal structure.

Metric Structure. The Riemannian metric ρ on S is defined by a differential element of arclength ds . Integrating ds along a path gives its length, and the distance $\rho(p, q)$ is the length of the shortest path lying in S and connecting p to q . Double integration of ds over a region $\Omega \subset S$ gives its area.

Curvature Structure. The *curvature* of S has to do with its shape in 3-space. We use *Gaussian curvature* since it depends only on distances in S itself, not on how S lies in \mathbb{R}^3 . A rounded region, say a hilltop, represents positive curvature (Figure 8(a)); a plain or valley floor is flat, zero curvature; while a saddle point, such as a mountain pass, has negative curvature (Figure 9(a)).

Since our surfaces are *piecewise flat* with triangular faces, the notion of curvature requires some interpretation. All the nonzero curvature resides at the vertices. If v is a vertex and T is the sum of the angles at v in all the triangles meeting at v , then

curvature at v is defined to be $2\pi - T$. Positive, zero, and negative curvature have their familiar geometric interpretations. For example, a cone point composed of 5 equilateral triangle faces has a total angle sum of $5\pi/3$, giving positive curvature $\pi/3$. Six faces of equilateral triangles demonstrates zero curvature (flat); and a saddle point composed with 8 equilateral faces demonstrates negative curvature $-2\pi/3$.

Conformal structure. Given two smooth oriented curves γ_1 and γ_2 in S which cross at a point p , the *angle* between γ_1 and γ_2 at p is the smallest turning angle in S which will bring the tangent vector of γ_1 at p to the tangent vector of γ_2 at p . A mapping $f : S_1 \mapsto S_2$ is said to be *conformal* if it preserves all such angles, both in magnitude and orientation.

On piecewise flat surfaces there is a slight subtlety about the meaning of angles between curves when they meet at a vertex v ; namely, one needs a “market share” or proportional interpretation. Suppose γ_1 and γ_2 intersect at v . As before, let T denote the sum of the angles of all the triangles meeting at v . Sweeping the curve γ_1 counterclockwise about v into γ_2 within the surface accumulates some turning angle β through these faces. The angle between γ_1 and γ_2 is not β , but rather is the ratio $\alpha = 2\pi\beta/T$ which measures β ’s *market share* or proportion of T . This is really quite natural: if β represents, say, a quarter of the total angle at v , $\beta = (\frac{1}{4})T$, then under a conformal flattening one would expect the images of γ_1 and γ_2 to meet in the angle $\alpha = (\frac{1}{4})2\pi = \pi/2$, that is, in a right angle.

These interpretations of curvature and angle in the piecewise flat setting are entirely standard. They are local, intrinsic to S , converge to the usual meanings when piecewise flat surfaces approximate smooth surfaces, and they have exactly the intuitive content one expects.

Primer on the Classical Geometries. The classical geometric spaces, \mathbb{S}^2 , \mathbb{C} , and \mathbb{D} , share several common features; let \mathcal{T} denote any one of the three spaces. Each \mathcal{T} enjoys a *Riemannian* geometry; this is associated with a metric ρ based on a differential element of arclength ds which gives lengths of curves and areas of regions as described earlier. A *geodesic*, or “straight line”, segment in \mathcal{T} is a shortest curve between two points, p and q , and the distance $\rho(p, q)$ is its length. A *circle* is the set of points a given ρ -distance r (the *radius*) from a fixed point p (the *center*). Gaussian curvature for each of our spaces can be computed directly from the metric in Table 1; in each case it is constant. Angles between curves are computed using a general version of the Law of Cosines which takes account of curvature.

Each space \mathcal{T} has a collection $\text{Aut}(\mathcal{T})$ of one-to-one conformal self-maps, $\phi : \mathcal{T} \rightarrow \mathcal{T}$, termed *automorphisms*. These are mathematical *groups* under composition, and their structures account for the practical advantages these spaces enjoy as targets for our flat maps. (For example, $\text{Aut}(\mathbb{S}^2) = \mathcal{M}$ is equivalent to $PSL(2, \mathbb{C})$, perhaps the most thoroughly studied group in mathematics and physics.) The automorphism groups share several useful properties: 1) Each $\phi \in \text{Aut}(\mathcal{T})$ maps circles to circles (indeed, in each case these are (up to orientation) the only self-maps of \mathcal{T} with this property). 2)

The orientation preserving *rigid* motions, or isometries, (automorphisms which also preserve distances) form a subgroup. 3) Each group $\text{Aut}(\mathcal{T})$ is *doubly transitive*, meaning roughly that given any points $p, q \in \mathcal{T}$ and any directions α, β , there exists $\phi \in \text{Aut}(\mathcal{T})$ (in fact, a rigid motion) so that $\phi(p) = q$ and a curve passing through p in direction α is mapped to a curve passing through q in direction β . In other words, all standard normalizations can be accomplished with automorphisms.

As pointed out earlier, our three classical geometries form a unified and nested hierarchy. A key ingredient is *stereographic* projection, which gives the inclusion $\mathbb{C} \subset \mathbb{S}^2$. Namely, every point z of \mathbb{C} is identified with the point p of \mathbb{S}^2 where the line from z to the south pole of \mathbb{S}^2 pierces \mathbb{S}^2 . See Figure 15(f). Stereographic projection preserves circles, angles, and orientation and respects automorphisms. It also allows for the transfer of complex arithmetic from \mathbb{C} to \mathbb{S}^2 , which we exploit in Table 1. Note that most texts describing stereographic projection puncture \mathbb{S}^2 at the north pole. If the south pole is used, the resulting stereographic map is more intuitive for the untrained user.

Here are observations about the individual geometries.

The Sphere. There is a simple geometric interpretation of the spherical metric ρ (see Table 1): given points $p, q \in \mathbb{S}^2$, view them from the origin; the angle between the two viewing directions (all angle measurements are in radians) is precisely the spherical distance $\rho(p, q)$. From this it's easy to deduce that a circle in \mathbb{S}^2 is the intersection of a Euclidean plane in \mathbb{R}^3 with \mathbb{S}^2 and that geodesic segments are arcs of great circles (circles of radius $\pi/2$). In particular, this is a non-Euclidean geometry: any two straight lines intersect. In normalizing our spherical maps we rely on the fact that for any two triples $\{p_1, p_2, p_3\}$ and $\{q_1, q_2, q_3\}$ of points of \mathbb{S}^2 there exists a unique $\phi \in \text{Aut}(\mathbb{S}^2)$ so that $\phi(p_j) = q_j, j = 1, 2, 3$.

The Plane. The distance between $z_1 = x_1 + iy_1$ and $z_2 = x_2 + iy_2$ is $|z_1 - z_2| = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$, the usual Pythagorean distance in \mathbb{R}^2 . Each $\phi \in \text{Aut}(\mathbb{C})$ is an *affine* map; that is, it is some combination of a rotation, a dilation, and/or a translation. The identification of \mathbb{C} under stereographic projection with \mathbb{S}^2 (minus the south pole) is conformal and maps circles of \mathbb{C} to circles of \mathbb{S}^2 , though it does not respect centers. Straight lines in \mathbb{C} are generally treated as circles “going through infinity” because they correspond under stereographic projection to circles in \mathbb{S}^2 starting and ending at the south pole.

The Disc. The automorphisms of \mathbb{D} preserve hyperbolic distances, and studying ds , one can fairly easily establish the properties we listed in Section 3 about circles and geodesics. For instance, suppose c is a Euclidean circle of radius $a < 1$. Center c at origin and its hyperbolic radius can be computed as the Poincarè metric for \mathbb{D} , $\rho(0, a) = \frac{1}{2} \log\left(\frac{1+a}{1-a}\right)$. Now move that Euclidean circle towards the boundary; ds tells us that its hyperbolic radius grows and that the hyperbolic center is displaced ever further towards the unit circle from its Euclidean center. The instant c touches $\partial\mathbb{D}$ at some point ξ , the center also hits ξ — this is a horocycle, infinite radius and ideal center at ξ . Hyperbolic geometry is a second type of non-Euclidean geometry: given a line L and a point q not on L , there are infinitely many lines through q parallel to L .

Primer on Surface Maps. A *map* (or *mapping*) f from one surface, S_1 , to another, S_2 , is a *one-to-one function* $f : S_1 \mapsto S_2$. That is, every point x_1 of S_1 corresponds with a unique point $x_2 = f(x_1)$ in S_2 , and if x and y are distinct points of S_1 then $f(x)$ and $f(y)$ are distinct points of S_2 . Thus f effectively identifies S_1 with the subset $f(S_1)$ in S_2 so that locations, curves, or regions of S_1 are identified with corresponding locations, curves, or regions in S_2 .

Assuming S_1 and S_2 have conformal structures, we've already noted that f is a *conformal map* if it preserves angles between curves. An equivalent and more intuitive condition may be paraphrased as follows: f is conformal if for each point p of S_1 it carries tiny circles centered at p to tiny (approximate) circles centered at $f(p)$ in S_2 . This latter condition can be loosened to define the more general class of *quasiconformal maps*: f is κ -*quasiconformal* if for each point p of S_1 it carries tiny circles centered at p to tiny (approximate) *ellipses* centered at $f(p)$ in S_2 and having eccentricity bounded by κ . For precise definitions see Lehto and Virtanen (1973).

There are innumerable equivalent (and precise) definitions of *quasiconformal*; it is a highly developed field often used in studies of conformal mapping. κ , which is greater than or equal to 1, is one measure of conformal *distortion* of κ -quasiconformal maps; indeed, f is conformal if and only if it is 1-quasiconformal. However, κ indicates “worst-case” local distortion and is not particularly relevant in practice. Even for large κ , a κ -quasiconformal map f can convey significant conformal information.

Primer on Circle Packing. As defined earlier, a circle packing is a configuration of circles with a specified pattern of tangencies. Here is some standard terminology:

Complex: K denotes a *complex* (technically, an abstract simplicial 2-complex) associated with a triangulation of a topological surface. This is basically a “list” of the vertices v , edges $\langle v, u \rangle$, and oriented faces $\langle v, u, w \rangle$ of the triangulation, and represents our “pattern”.

Packing: P denotes a *circle packing* for K in one of our target geometries \mathcal{T} . More specifically, P is a collection $\{c_v\}$ of circles, one for each vertex v of K , so that c_v is tangent to c_u whenever $\langle v, u \rangle$ is an edge of K and so that the triple $\langle c_v, c_u, c_w \rangle$ of mutually tangent circles is positively oriented in \mathcal{T} whenever $\langle v, u, w \rangle$ is a positively oriented face of K .

Packing Label: R is the collection of radii associated with P ; for each vertex v of K , r_v is the radius (in the metric of \mathcal{T}) for the circle c_v .

Carrier: $\text{carr}(P)$ denotes the concrete *geometric triangulation* in \mathcal{T} formed by connecting the centers of tangent circles of P with geodesic segments; this provides a mesh in \mathcal{T} which is combinatorially equivalent to K .

In our circumstances: 1) The complex K is the triangular mesh of some reconstructed cortical surface S . 2) The packing P is a circle packing for K guaranteed by the Discrete Riemann Mapping Theorem or the Circle Packing Theorem (Beardon and Stephenson, 1990). 3) The packing label R is what one computes (approximately) so that P can be laid out. 4) $\text{Carr}(P)$ is used (as described below) to define the desired

flat map $f : S \longrightarrow \mathcal{T}$. Recalling that we are starting with a given topologically correct complex K , here are some of the technical issues involved:

A) The computational effort lies in approximating packing labels R . With R in hand, it is straightforward to construct P itself. Starting (in the appropriate geometric space \mathcal{T}) with one circle and a tangent neighbor, one can successively lay out additional circles which have two contiguous neighbors already in place. The packing condition and respect for orientation ensure that there will be no ambiguity in placement. The outcome is determined entirely by the first two circles, giving the three degrees of freedom used in normalizations: coordinates of the first circle’s center and an angle for the second circle’s point of tangency.

B) There is (as yet) no packing algorithm intrinsic to spherical geometry. One uses the puncture trick described in Section 5 so the computations are actually done in \mathbb{D} . Other triangulated surfaces can also be packed in \mathbb{S}^2 , as illustrated in Figure 15(e), but all are projected from \mathbb{C} or \mathbb{D} .

C) Suppose P is a circle packing for K in the target space \mathcal{T} . Each vertex p of S corresponds to a vertex v of K and hence to a circle c_v of P . Define $f(p) = z_v$, where z_v is the center of c_v in \mathcal{T} . In this way the three vertices of any face of S are identified with the three vertices of the corresponding triangle in $\text{carr}(P)$. We complete the definition of f by extending, first to the edges of S and then to the faces. *The map $f : S \mapsto \text{carr}(P) \subset \mathbb{R}^2$ is the “flat” map on S induced by this particular packing P , called a discrete conformal map.* Of course, other packings P for K will induce alternate maps. The main properties of f depend heavily on the geometry in which P lives, but details in the definition of f , such as how to extend it from vertices to edges and faces, are largely immaterial in practice.

D) A discrete conformal map $f : S \longrightarrow \mathcal{T}$ is κ -quasiconformal. One can establish an upper bound on the dilatation κ based on two items of data. First is the *degree* of K , the maximal number of edges emanating from any one vertex; the bound here derives from the Ring Lemma of Rodin and Sullivan (1987). Second is the range of angles in the faces of the triangulation of the surface S ; the closer to $\pi/3$ (60 degrees), the better this part of the bound. In practice, precise bounds for κ are not a high priority. As mentioned earlier, κ is worst-case — generic behavior appears to have much less conformal distortion. Moreover, our experience suggests that many of the advantages of the discrete theory are intrinsic and do not depend on mere approximation of the classical maps. In any case, mesh refinements, new computational approaches based on *inversive distance* (a hyperbolic invariant) packings, and parallelization promise continual improvements in these approximations.

E) All computations, manipulations, and visualizations of circle packings displayed in this paper are implemented in the software suite **CirclePack**. The visual presentation of a circle packing flat map uses either the circles of the packing or the faces of carrier. Any combination of circles and/or faces, filled/unfilled, color/b&w, solid/semitransparent can be chosen at any time. A color index is attached to each vertex or face for color

rendering. `CirclePack` floods the interior of the object (circle or face) with the specified color and renders its boundary in a slightly more saturated shade of the same color, as this seems to give the viewer the best appreciation for the structure underlying the flat map. Viewing of 3D surfaces requires a separate graphical interface. Many packages are available for this and can be linked to `CirclePack`. Points or regions chosen on one surface can be highlighted on another. Möbius transformations, which allow the hyperbolic map focus or the spherical map poles and equator to be changed interactively, occur in real time.

REFERENCES

- Ahlfors, L. V. 1966. *Complex Analysis*. McGraw-Hill Book Company, New York.
- Andre'ev, E. M. 1970. Convex polyhedra in Lobacevskii space. *Math. USSR Sbornik* **10**:413–440.
- Armstrong, M. A. 1983. *Basic Topology*. Springer-Verlag, New York.
- Beardon, A. F. and Stephenson, K. 1990. The uniformization theorem for circle packings. *Indiana Univ. Math. J.* **39**:1383–1425.
- Collins, C. and Stephenson, K. 2001. A circle packing algorithm. Submitted. Available at <http://www.math.utk.edu/~kens>.
- Dale, A. M. and Sereno, M. I. 1993. Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: A linear approach. *Journal of Cognitive Neuroscience* **5**:162–176.
- Dale, A. M., Fischl, B., *et al.* 1999. Cortical surface-based analysis I: Segmentation and surface reconstruction. *Neuroimage* **9**:179–194.
- Divernoy, H. M. 1995. *The Human Brain Stem and Cerebellum*. Springer-Verlag, Wein.
- Drury, H. A. and Van Essen, D. C. 1997. Functional specializations in human cerebral cortex analyzed using the visible man surface-based atlas. *Human Brain Mapping* **5**:233–237.
- Drury, H. A., Van Essen, D. C., *et al.* 1996. Computerized mappings of the cerebral cortex: A multiresolution flattening method and a surface-based coordinate system. *Journal of Cognitive Neuroscience* **8**:1–28.
- Fischer, B. 1973. Overlap of receptive field centres and representation of the visual field in the cat's optic tract. *Vision Research* **13**:2113–2120.
- Fischl, B., Sereno, M. I., *et al.* 1999. Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. *Neuroimage* **9**:179–194.
- Holmes, C. J., Hoge, R., *et al.* 1996. Enhancement of T1 MR images using registration for signal averaging. *The Journal of Neuroscience* **3**:S28. Part 2 of 2.
- Koebe 1936. Kontaktprobleme der Konformen Abbildung. *Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl.* **88**:141–164.
- Lehto, O. and Virtanen, K. I. 1973. *Quasiconformal Mappings in the Plane*. Second edn. Springer-Verlag, New York.
- Lorensen, W. E. and Cline, H. 1987. Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics* **21**:163–169.
- Massey, W. S. 1967. *Algebraic Topology: An Introduction*. Springer-Verlag, New York.
- Murray, J. D. 1989. *Mathematical Biology*. Springer-Verlag, Berlin, chapter 16.
- Polya, G. 1968. *Mathematical Discovery*. Vol. 2. John Wiley & Sons, New York.
- Rehm, K., Lakshminaryan, K., *et al.* 1998. A symbolic environment for visualizing activated foci in functional neuroimaging datasets. *Medical Image Analysis* **2**:215–226.
- Rehm, K., Rottenberg, D., *et al.* 2000. Use of cerebellar landmarks to define a coordinate system and an isolation strategy. *Neuroimage* **11**:S536. Part 2.
- Rodin, B. and Sullivan, D. 1987. The convergence of circle packings to the Riemann mapping. *J. Differential Geometry* **26**:349–360.

- Rosa, M. G. P., Fritsches, K. A., *et al.* 1997. The second visual area in the marmoset monkey: Visuotopic organisation, magnification factors, architectonical boundaries, and modularity. *The Journal of Comparative Neurology* **387**:547–567.
- Schmahmann, J. D., Doyon, J., *et al.* 1999. Three-dimensional MRI atlas of the human cerebellum in proportional stereotaxic space. *Neuroimage* **10**:233–260.
- Schroeder, W., Martin, K., *et al.* 1998. *The Visualization Toolkit*. Prentice Hall, Upper Saddle River, NJ.
- Schwartz, E. L. 1980. Computational anatomy and functional architecture of striate cortex: A spatial mapping approach to perceptual coding. *Vision Research* **20**:645–669.
- Schwartz, E. L. 1994. Computational studies of the spatial architecture of primate visual cortex. *Primary Visual Cortex in Primates* (A. Peters and K. S. Rockland, Eds.). Vol. 10 of *Cerebral Cortex*, chapter 9, pp. 375–395. Plenum Press, New York. General editor J. R. Cronly-Dillon.
- The National Library of Medicine 2000. The visible human project. Web document available: <http://www.nlm.nih.gov/research/visible/>.
- Thurston, W. 1985. The finite Riemann mapping theorem. Invited talk, An International Symposium at Purdue University on the occasion of the proof of the Bieberbach conjecture, March 1985.
- Thurston, W. 1997. *The Geometry and Topology of 3-Manifolds*. Princeton University Press, Princeton, NJ.
- Toga, A. W. and Mazziotta, J. C. (Eds.) 1996. *Brain Mapping: The Methods*. Academic Press, Inc., San Diego.
- Tusa, R. J., Palmer, L. A., *et al.* 1978. The retinotopic organization of area 17 (striate cortex). *Journal of Comparative Neurology* **177**:213–236.
- Van Essen, D. C. and Drury, H. A. 1997. Structural and functional analyses of human cerebral cortex using a surface-based atlas. *The Journal of Neuroscience* **17**:7079–7102.
- Van Essen, D. C., Drury, H. A., *et al.* 1998. Functional and structural mapping of human cerebral cortex: Solutions are in the surfaces. *Proceedings of the National Academy of Sciences, U.S.A.* **95**:788–795.

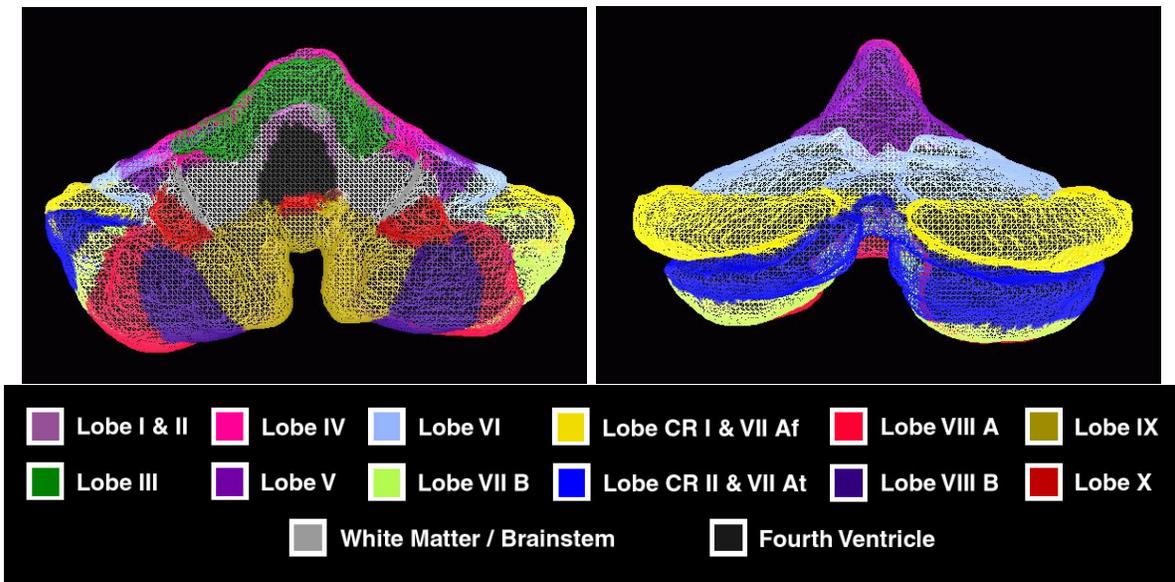


FIGURE 1. Anterior and posterior views of triangulated surface of the human cerebellum (56,676 triangles; 28,340 vertices). For the planar flat maps, a boundary corresponding to the white matter cut-plane and filled-in fourth ventricle was introduced where the brainstem attaches to the cerebellum.

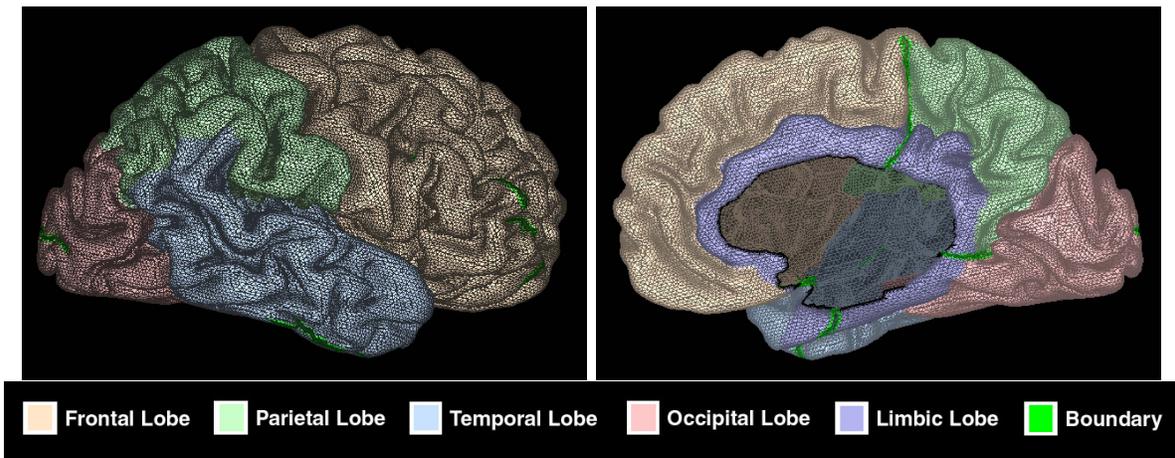
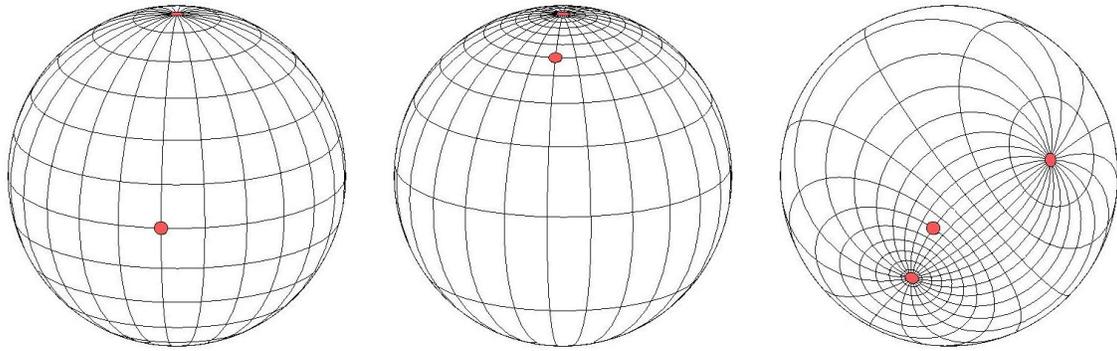


FIGURE 2. Lateral and medial views of triangulated surface of the Visible Man right cerebral hemisphere (103,845 triangles; 52,360 vertices).



(a) Spherical Map

(b) Möbius transformation that shifts E.

(c) Möbius transformation that shifts N, S, E.

FIGURE 3. Automorphisms of the sphere \mathbb{S}^2 . North (N) and south (S) poles and an equatorial point (E) are marked in red.

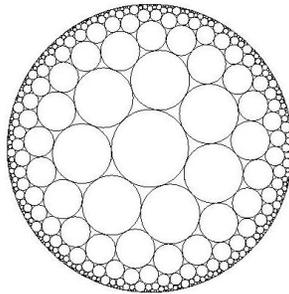
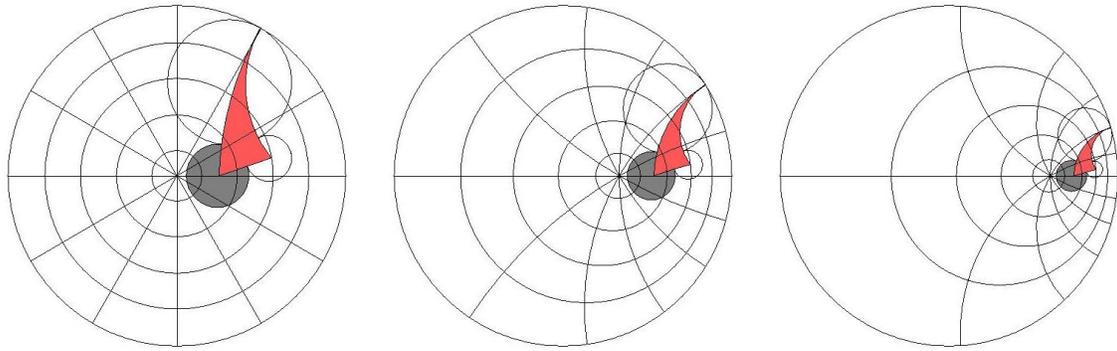


FIGURE 4. Each circle is the same hyperbolic area. To our Euclidean eyes, the circles appear to be shrinking, but in the hyperbolic plane they are located farther and farther away. As a circle approaches the boundary it approaches an infinite distance from the origin.



(a) Hyperbolic map.

(b) Möbius transformation applied.

(c) Möbius transformation applied again.

FIGURE 5. Automorphisms of the hyperbolic plane \mathbb{D} . The automorphism $\phi : z \mapsto (3z + 1)/(3 + z)$ is applied successively. In practice, this allows selected regions to be brought into focus while relegating other regions to the periphery of the hyperbolic map.

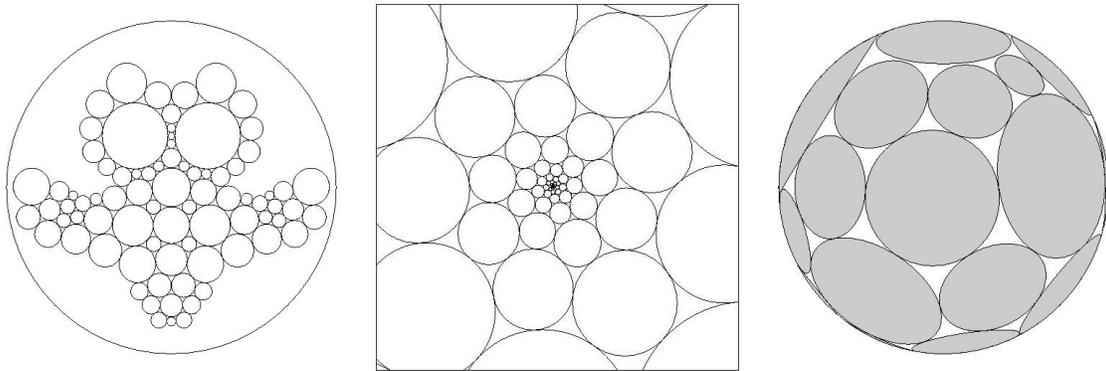


FIGURE 6. Examples of circle packings.

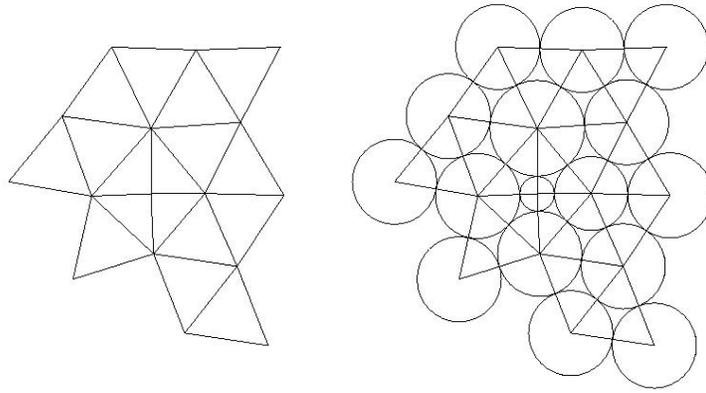


FIGURE 7. A triangulation and its corresponding circle packing.

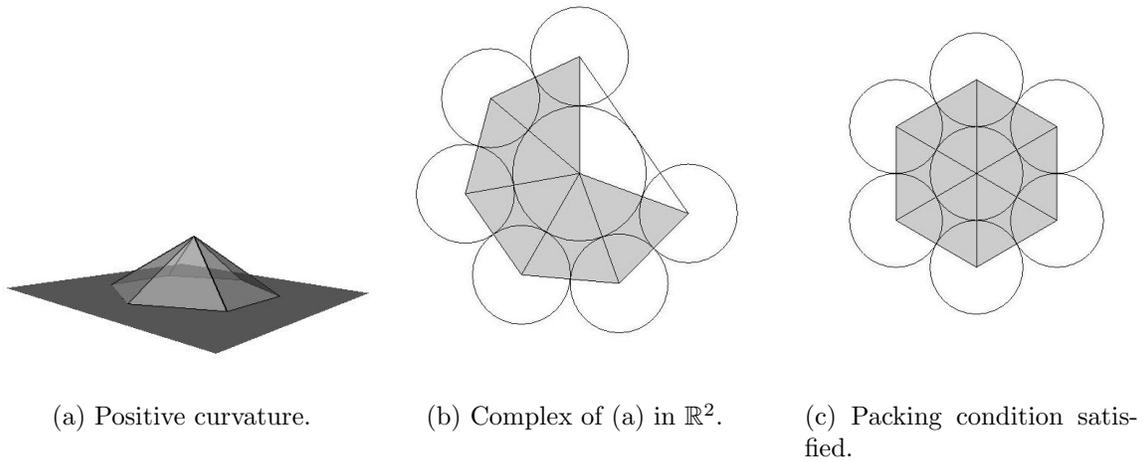


FIGURE 8. Obtaining a circle packing when there is positive curvature. Since the angle sum is less than 2π , the radius of the interior circle must be decreased to satisfy the packing condition and make all circles tangent.

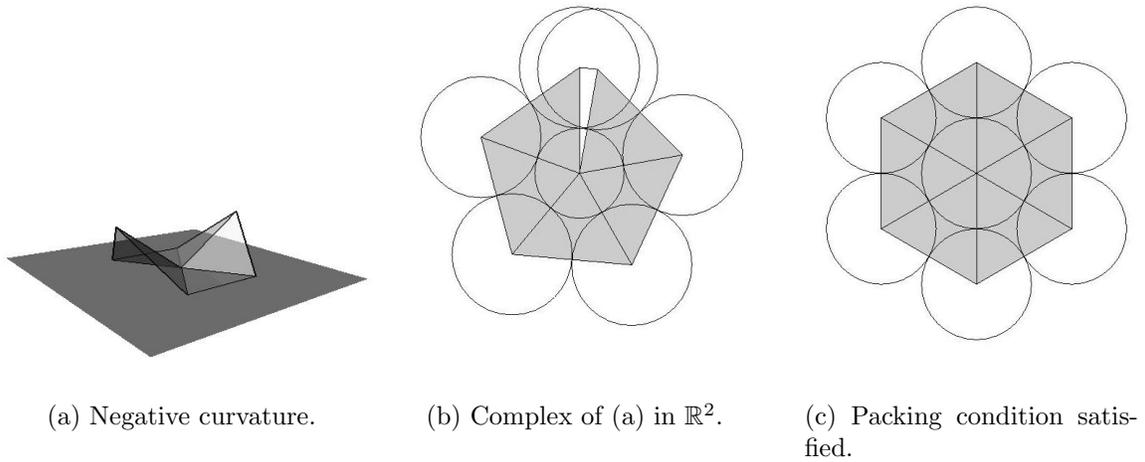


FIGURE 9. Obtaining a circle packing when there is negative curvature. Since the angle sum is greater than 2π , the radius of the interior circle must be increased to satisfy the packing condition and make all circles tangent.

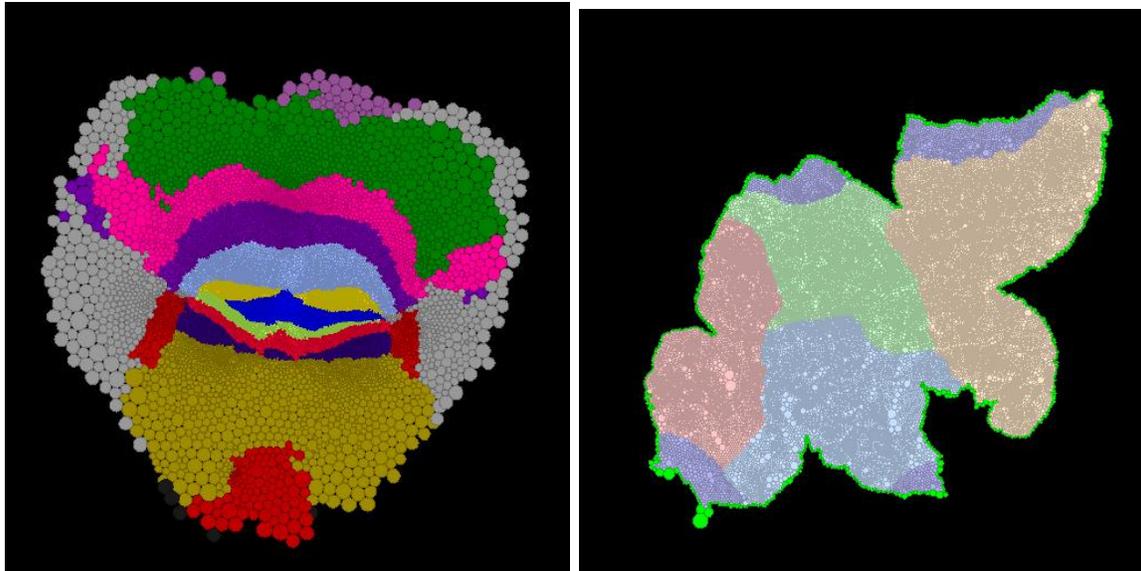


FIGURE 10. Euclidean flat maps of the cerebellum (left) and cerebrum (right). Colors are those of Figures 1 and 2.

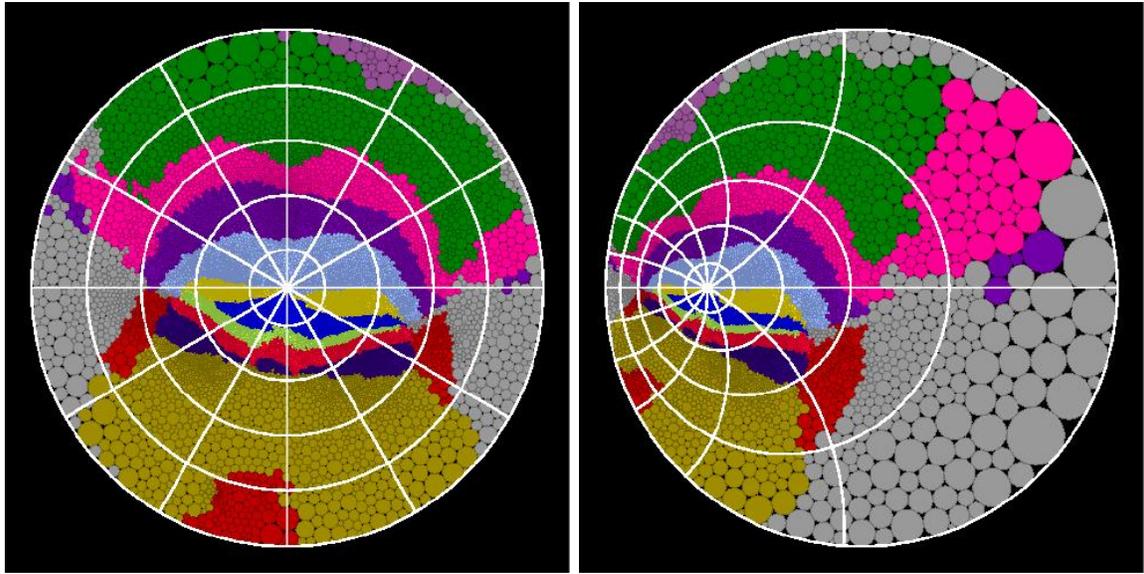


FIGURE 11. Hyperbolic flat maps of the cerebellum with different Möbius transformations applied. A polar Euclidean canonical grid system has been imposed (compare with grid of Figure 14). Colors are those of Figure 1.

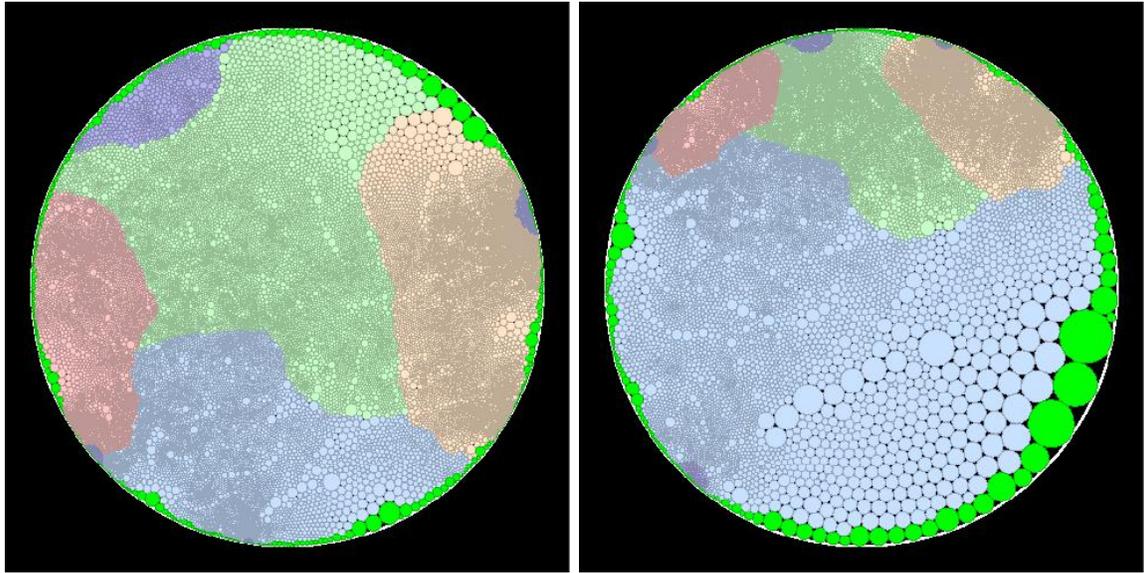
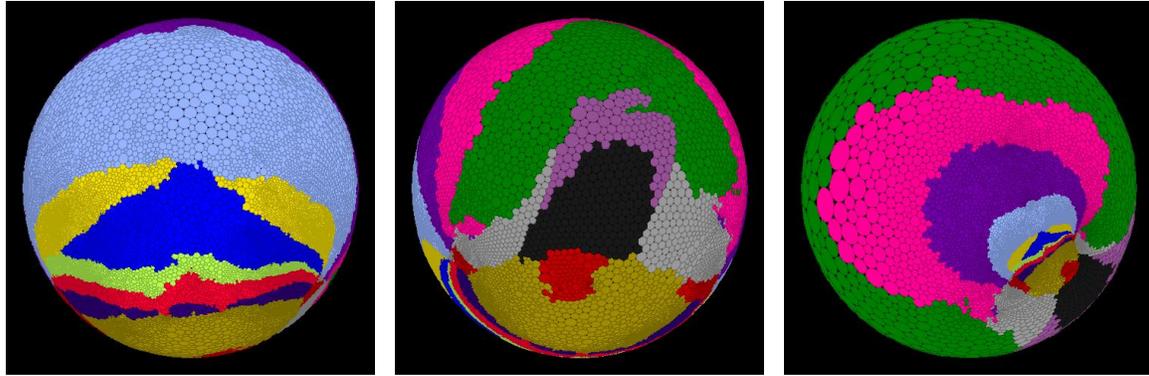


FIGURE 12. Hyperbolic flat maps of the cerebrum with different Möbius transformations applied. Colors are those of Figure 2.

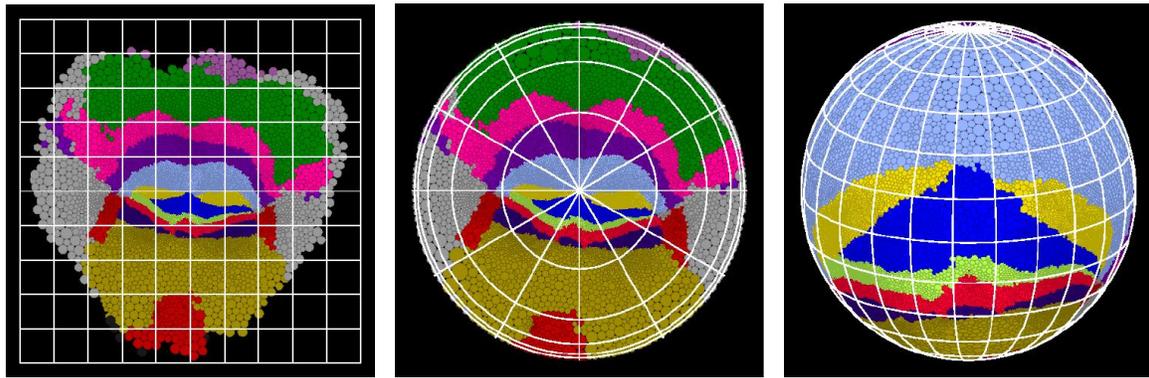


(a) Spherical map.

(b) Map rotated.

(c) Application of an automorphism.

FIGURE 13. Spherical flat maps of the cerebellum. Colors are those of Figure 1.

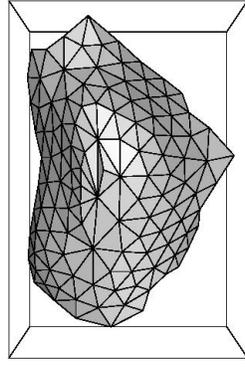
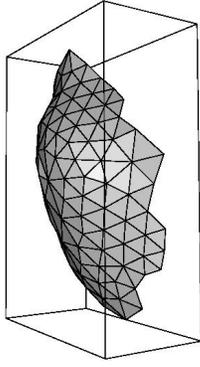


(a) Euclidean map.

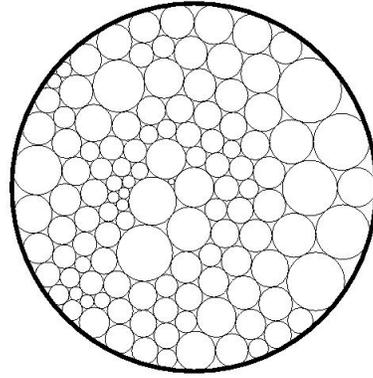
(b) Hyperbolic map.

(c) Spherical map.

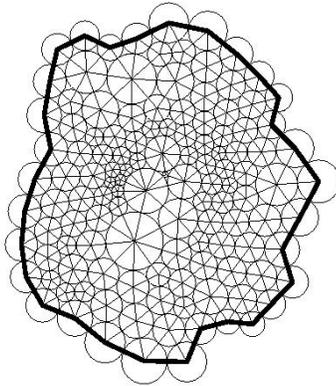
FIGURE 14. Flat maps of the cerebellum with canonical coordinate grid systems imposed. The local structures in the interior of the maps are nearly identical. The hyperbolic map uses a grid of hyperbolic distances rather than Euclidean distances (compare with grid of Figure 11). Colors are those of Figure 1.



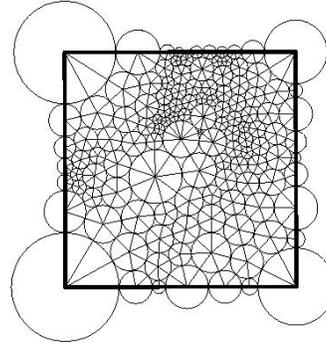
(a) Subsurface of Figure 2.



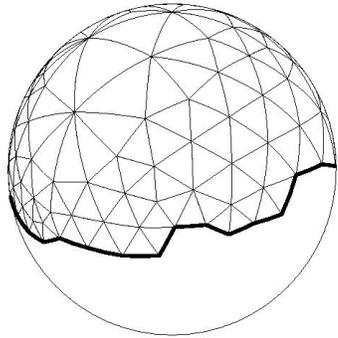
(b) Packing on the hyperbolic disk.



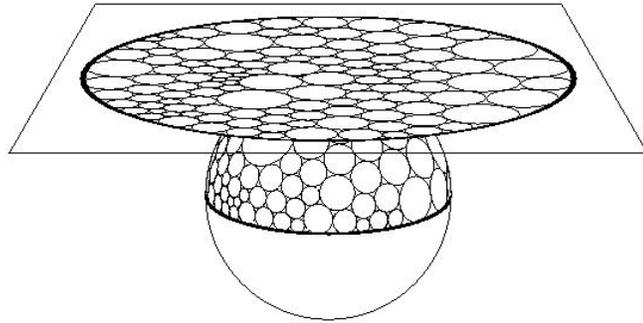
(c) Packing specifying boundary lengths.



(d) Packing in the plane into a rectangle.



(e) Packing on the sphere.



(f) Stereographic projection onto the hemisphere.

FIGURE 15. Quasi-conformal maps of the same complex using different labels. Complex is a sub-surface of Figure 2 (211 triangles; 122 vertices including 31 boundary vertices/edges).

TABLE 1. Properties of the Three Different Geometries.

| Property | Description | Geometry |
|---|---|------------|
| Geometric model | sphere $\mathbb{S}^2 = \{(x, y, z) : x^2 + y^2 + z^2 = 1\}$ | Spherical |
| | Euclidean plane \mathbb{R}^2 or complex plane \mathbb{C} | Euclidean |
| | unit disc $\mathbb{D} = \{(x, y) : x^2 + y^2 < 1\}$ | Hyperbolic |
| Points denoted by | $P = (x, y, z)$ | Spherical |
| | $P = (x, y)$ in \mathbb{R}^2 or $z = x + iy$ in \mathbb{C} | Euclidean |
| | $z = x + iy$ | Hyperbolic |
| Differential element of arclength | $ds = \sqrt{dx^2 + dy^2 + dz^2}$ | Spherical |
| | $ds = \sqrt{dx^2 + dy^2}$ | Euclidean |
| | $ds = \frac{2 dz }{1 - z ^2}$ | Hyperbolic |
| Metric ρ (distance between points) | $\rho(P_j, P_k) = \arccos(P_j \cdot P_k)$ $= \arccos(x_j x_k + y_j y_k + z_j z_k)$ (radians) | Spherical |
| | $\rho(P_j, P_k) = \sqrt{(x_j - x_k)^2 + (y_j - y_k)^2}$ | Euclidean |
| | $\rho(z_j, z_k) = \frac{1}{2} \log \left(\frac{ 1 - z_j \bar{z}_k + z_j - z_k }{ 1 - z_j \bar{z}_k - z_j - z_k } \right)$ | Hyperbolic |
| Curvature and geodesics | curvature +1; geodesics are great circles | Spherical |
| | curvature 0; geodesics are straight lines | Euclidean |
| | curvature -1; geodesics are circle arcs that intersect the boundary of \mathbb{D} orthogonally | Hyperbolic |
| Conformal automorphisms | $\phi : z \mapsto \frac{az + b}{cz + d}$, $a, b, c, d \in \mathbb{C}$ with $ad - bc \neq 0$ and where z is projected stereographically (full Möbius group) | Spherical |
| | $\phi : z \mapsto az + b$, $a, b \in \mathbb{C}$, $a \neq 0$ | Euclidean |
| | $\phi : z \mapsto e^{i\theta} \left(\frac{z - \alpha}{1 - \bar{\alpha}z} \right)$, $\theta \in \mathbb{R}$, $\alpha \in \mathbb{D}$ | Hyperbolic |
| Packing condition for label $R = \{r_v\}$ | $\sum_{\langle v, u, w \rangle} \arccos \left\{ \frac{\cos(r_u + r_w) - \cos(r_v + r_u) \cos(r_v + r_w)}{\sin(r_v + r_u) \sin(r_v + r_w)} \right\} = 2\pi$ | Spherical |
| | $\sum_{\langle v, u, w \rangle} \arccos \left\{ \frac{(r_v + r_u)^2 + (r_v + r_w)^2 - (r_u + r_w)^2}{2(r_v + r_u)(r_v + r_w)} \right\} = 2\pi$ | Euclidean |
| | $\sum_{\langle v, u, w \rangle} \arccos \left\{ \frac{\cosh(r_v + r_u) \cosh(r_v + r_w) - \cosh(r_u + r_w)}{\sinh(r_v + r_u) \sinh(r_v + r_w)} \right\} = 2\pi$ | Hyperbolic |