

# Discrete Conformal Mappings via Circle Patterns

Liliya Kharevych  
Caltech

Boris Springborn  
TU Berlin

Peter Schröder  
Caltech



Figure 1: Examples of discrete conformal maps produced with our method. Next to each 3D texture image is a visualization of the planar region over which the surface is parameterized. The (cut) Camel demonstrates a constrained complex boundary shape; the Max Planck parameterization (also with cuts) shows a straightline bounded parameter region suitable for good texture packing; the face mask demonstrates natural boundary conditions; and the lion head mapping to a disk.

## Abstract

We introduce a novel method for the construction of discrete conformal mappings from (regions of) embedded meshes to the plane. Our approach is based on *circle patterns*, i.e., arrangements of circles—one for each face—with prescribed intersection angles. Given these angles the circle radii follow as the *unique* minimizer of a convex energy. The method has two principal advantages over earlier approaches based on discrete harmonic mappings: (1) it supports very flexible boundary conditions ranging from natural boundaries to control of the boundary shape via prescribed curvatures; (2) the solution is based on a convex energy as a function of *logarithmic* radius variables with simple explicit expressions for gradients and Hessians, greatly facilitating robust and efficient numerical treatment. We demonstrate the versatility and performance of our algorithm with a variety of examples.

**CR Categories:** G.1.0 [Numerical Analysis]: General—Numerical Algorithms; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid and object representations; Geometric algorithms, languages, and systems.

**Keywords:** Conformal parameterizations; discrete differential geometry; circle patterns; discrete analytic functions; meshing; texture mapping

## 1 Introduction

Surfaces are often represented as collections of samples with connectivity, typically in the form of a simplicial mesh. It is natural

and convenient to use the implied piecewise linear mesh as the basis for formulating a variety of computational algorithms, such as parameterization problems or the solution of partial differential equations for purposes of simulation. In this paper we argue that when it comes to computing conformal structures, e.g., conformal parameterizations of surfaces, circles can be a far better basis upon which to formulate the underlying relationships and consequent algorithms (see Figure 2). In particular we advocate the formulation of the *discrete conformal mapping*<sup>1</sup> problem in terms of circles and the angles with which they intersect, so called *circle patterns*.

The idea of using circles to capture a discrete notion of conformality goes back to a conjecture of Thurston’s [1985] who posited that one may approximate the Riemann mapping<sup>2</sup> from a given region in the plane to the unit disk through a sequence of increasingly fine, regular (hexagonal) *circle packings*. This conjecture was later proven correct by Rodin and Sullivan [1987]. Circle packings assign a circle to each vertex, with pairwise tangency for each edge in the mesh. A numerical algorithm for the construction of such mappings, based on iterative adjustment of circle radii, was proposed by Thurston [1980] and improved and realized by Collins and Stephenson [2003]. Unfortunately, circle *packings* yield mappings which depend only on the *combinatorics* of the original mesh, while we are seeking methods which depend on the *geometry* of the mesh. One possible avenue to remedy this shortcoming is to use patterns of non-intersecting circles [Bowers and Hurdal 2003]. Unfortunately there exists little theory concerning these patterns and

<sup>1</sup>In this article we use the term “discrete conformal map” for maps between meshes that are close to angle preserving.

<sup>2</sup>The Riemann Mapping Theorem asserts that there exists a unique (up to Möbius transformations of the unit disk to itself) conformal map from any region in the plane (open, connected and simply connected, not the whole plane) onto the unit disk.

in empirical practice solutions cannot always be found.

In contrast, circle *patterns*, which associate a circle with each *face* in the original mesh provide an opportunity to incorporate the intrinsic geometry of the original mesh: each edge is assigned an angle  $\theta \in (0, \pi)$  which corresponds to the intersection angle of the two incident face circles. A recent theorem of Bobenko and Springborn [2004] characterizes such circle patterns as the *unique* minimizer of a convex energy expressed in terms of *logarithmic* radius variables (and the given edge angles). Simple, explicit expressions for the energy, its gradient, and Hessian are available and greatly facilitate an efficient implementation.




Figure 2: Typically a triangle mesh is understood as the piecewise linear interpolation of given vertex coordinates induced by the connectivity of the mesh (left). Alternatively we may also think of the vertices as the unique loci where incident triangle circumcircles intersect (middle;right). The latter point of view is more appropriate for formulating relationships of conformal geometry.

**Contributions** We describe a robust and efficient numerical procedure for the construction of discrete conformal mappings. In the first stage of the algorithm, we compute edge labels that are close to the circumcircle intersection angles in the original mesh and may serve as circumcircle intersection angles of a planar Delaunay mesh. This requires solving a quadratic programming problem. In the second stage, we use the theory of Bobenko and Springborn [2004] to construct the planar Delaunay mesh with the given intersection angles from the first stage. Here we have to minimize a convex function subject to *no constraints*. The variables are the logarithmic circumcircle radii. We modify the energy of Bobenko and Springborn (without compromising the underlying theory) to provide a simple and uniform treatment of boundary conditions (see Figure 3). With these we can provide both natural boundaries and detailed control over the boundary shape (see Figure 1). The parameter mappings are always locally injective. They may fail to be globally injective due to self-overlap of the boundary of the parameter domain. However, this can be avoided since we can prescribe the boundary curvature  $\kappa$ : if for any sequence of consecutive boundary vertices the sum of  $\kappa$ s is larger or equal  $-\pi$ , then there can be no overlap and the method is guaranteed to produce a global embedding. The resulting flat triangulations are of high quality as they are always triangulations of a (not necessarily convex) planar domain with all interior edges satisfying the local Delaunay criterion.




Figure 3: A simple mesh (left) mapped to the plane using circle patterns. Boundary shape is controlled through appropriate curvature conditions. Examples: disk boundary; natural boundary; and rectangular boundary (left to right).

## 1.1 Related Work

Most approaches to the construction of conformal mappings for meshes have relied on *discretizations of continuous formulations*. First order finite difference approximations of the Cauchy-Riemann equations were used by Levy *et al.* [2002]. The same equations, the so-called “cotan formula” [Pinkall and Polthier 1993], also result when considering a discrete variational Ansatz based on mesh invariants [Desbrun *et al.* 2002] or when deriving discrete holomorphy [Mercat 2001] from first principles using Discrete Exterior Calculus (DEC) methods [Hirani 2003]. Such approaches, for example, have been used to compute discrete approximations of Riemann structures for general meshes by Gu and co-workers (see [Gu and Yau 2003] and the references therein).

A great advantage of these methods is that they require only the solution of simple linear systems. However, due to negative cotan weights solutions may lack local injectivity. More troublesome is the lack of flexible boundary conditions: Dirichlet conditions introduce non-conformal distortion (see Figure 4), while natural boundaries provide essentially no control over the boundary.



Figure 4: Harmonic parameterization (left) of a region (middle) to the disk with Dirichlet boundary conditions (vertices are mapped to a  $k$ -gon with matching secant edge lengths) compared to our approach which sets appropriate angle conditions at the boundary (right).

A completely different Ansatz to the construction of conformal maps is based on circle *packing*. Continuous conformal mappings can be characterized as mapping infinitesimal circles to infinitesimal circles. Circle packings replace infinitesimal circles with finite circles. In the limit of refinement the continuous conformal maps are recovered [Rodin and Sullivan 1987]. Collins and Stephenson [2003] have implemented these ideas in their software *CirclePack*. The disadvantage of using circle packings (with tangent circles) is that they depend only on the combinatorics of the original mesh. In particular, if one starts with a planar mesh and parameterizes it, the result is not the original mesh. This is in contrast to our approach. Given a triangulation of a region in the plane satisfying the empty circumcircle property, one simply assigns the observed intersection angles of circumcircles (at the boundary one adds infinite circles, *i.e.*, straight lines) to each edge and our variational approach will return the identity map as a conformal map of the region to itself.

An extension of Stephenson’s original circle packing scheme that takes the geometry of the original mesh into account is based on patterns of non-intersecting circles, so called *inversive distance circle packings* [Bowers and Hurdal 2003]. (Non-intersecting circles have imaginary “intersection angles.”) However, virtually nothing is known regarding the existence and uniqueness of inversive distance packings.

The first variational principle for circle packings was presented in a seminal paper by Colin de Verdière [1991]. The variables are the circle radii. However, a closed formula is presented only for the derivative of the energy, not for the energy itself. Since then, different variational principles for circle packings [Brägger 1992] as well as circle patterns [Rivin 1994], [Leibon 2002] were discovered. In these, the variables are the angles of a triangulation,

subject to numerous linear constraints (one per edge and one per face for Rivin’s energy. Leibon’s energy deals with patterns in the hyperbolic plane.) In this paper, we use the most general functional which was given by Bobenko and Springborn [2004]. In their setup, the variables are logarithmic circle radii. Most importantly, they are not subject to any constraints.

Most closely related to our approach is the work of Sheffer and de Sturler [2000]. They flatten a given mesh by formulating a constrained quadratic minimization problem which seeks to find angles at the corners of triangles which are close to desired angles in a weighted  $L_2$  norm. The constraints capture the angle sum conditions at all faces (sum of angles  $= \pi$ ) and vertices (sum of angles  $= 2\pi$ ) as well as a non-linear condition on the product of sines of angles. The resulting minimization problem has local minima and does not have a unique solution. The resulting mappings are of excellent quality but only natural boundary conditions are provided. Their approach is similar in spirit to ours as we also optimize angles. We discuss the similarities and differences of both schemes in Section 2.3.

## 2 Circle Patterns and Discrete Conformal Maps

For computational problems in Euclidean geometry the use of triangles as a basic primitive is convenient and natural, since triangles are the basic invariant “building blocks” of Euclidean geometry. When one is interested in *conformal* geometry the picture changes. The basic invariants of conformal geometry are circles and the angles they make with one another. In the case of triangle meshes these differing points of view are naturally compatible. For example, the vertices in a triangle mesh are the unique loci where the circumcircles of the incident triangles intersect (see Figure 2). Similarly the empty circumcircle property, which corresponds to non-negative intersection angles between circumcircles incident on an edge, is a defining feature of Delaunay triangulations. While we may use triangles for tasks such as interpolation and rendering, conformal relationships between vertices are better captured through expressions involving the circumcircles they define and the angles these circles make with one another. A benefit of this different point of view is that much mathematical machinery from conformal geometry carries over to the discrete computational setting. For example, existence and uniqueness properties of conformal maps are reflected in the existence and uniqueness properties of circle patterns.

We begin this section by defining circle patterns in the plane and describe their characterization as minimizers of a variational energy. The latter forms the basis for our approach. While we only deal with triangle meshes here, the theory extends to polyhedral meshes and also to circle patterns with cone singularities [Bobenko and Springborn 2004].

### 2.1 Circle Patterns

Consider a Delaunay triangulation  $\mathcal{T} = (V, E, T)$  of finitely many points  $P = \{p_i\}$  in the plane. Here  $V = \{v_i\}$ ,  $E = \{e_{ij}\}$ , and  $T = \{t_{ijk}\}$  denote the sets of vertices, edges, and triangles respectively with  $p_i$  the point position of vertex  $v_i$ . From this Delaunay triangulation, we may read off the following edge weights

$$\forall e_{ij} \in E : \theta_e = \begin{cases} \pi - \alpha_{ij}^k - \alpha_{ij}^l & \text{for interior edges} \\ \pi - \alpha_{ij}^k & \text{for boundary edges} \end{cases}, \quad (1)$$

where  $\alpha_{ij}^k$  (and  $\alpha_{ij}^l$ ) are the angle(s) opposite  $e_{ij}$  in the adjacent triangle(s)  $t_{ijk}$  (and  $t_{jil}$ ). The  $\theta$ -weight of an interior edge is the (exterior) intersection angle of the circumcircles of the incident triangles (see Figure 5). The  $\theta$ -weight of a boundary edge is the intersection




Figure 5: Notation: The angles  $\alpha_{ij}^k$  and  $\alpha_{ij}^l$  opposite a given edge  $e_{ij}$  and incident on a particular vertex  $v_k$  respectively  $v_l$ . The edge angle  $\theta_e$  denotes the exterior intersection angle of the incident circumcircles or equivalently the angle between the two radii at  $v_i$  (or  $v_j$ ).

angle of the circumcircle of the incident triangle with the straight line containing the boundary edge. (View this line as a circle with infinite radius.)

Assume that the Delaunay triangulation is unique (points are in general position). Then

$$\forall e_{ij} \in E : 0 < \theta_e < \pi. \quad (2)$$

(Otherwise, some  $\theta_e$  may be 0.) For an interior vertex  $v_i$ , the sum of edge weights on the incident edges is  $2\pi$ :

$$\forall v_i \in V_{\text{int}} : \sum_{e \ni v_i} \theta_e = 2\pi, \quad (3)$$

while for a boundary vertex  $v_i$ , the defect

$$\forall v_i \in V_{\text{bdy}} : \kappa_i = 2\pi - \sum_{e \ni v_i} \theta_e \quad (4)$$

is the curvature angle of the polygonal boundary at that vertex (see Figure 6). Since the Delaunay triangulation triangulates the con-




Figure 6: The  $\theta$  angle sum at the boundary contains the discrete curvature term  $\kappa$ .

vex hull of the sites,  $0 \leq \kappa_i < \pi$ . However, we want to consider a slightly more general setup: *Instead of a Delaunay triangulation, we may start with a flat PL-surface that is topologically a disk and that is triangulated in such a way that the edge weights  $\theta$  satisfy the local Delaunay condition (Equation 2).* That is, we allow “Delaunay triangulations” of non-convex regions with polygonal boundary. Hence,  $\kappa_i$  may be negative, and we speak of *circle patterns* instead of Delaunay triangulations.

Now the idea is to reconstruct a circle pattern from its abstract triangulation and the intersection angles.

**Circle Pattern Problem** Given an abstract triangulation  $\mathcal{T}$  of a topological disk and a function  $\theta \in \mathbb{R}^E$  on the edge set  $E$  that satisfies Equation 2 and the angle sum condition for interior vertices (Equation 3), find a circle pattern that is combinatorially equivalent to  $\mathcal{T}$  and has the given edge weights  $\theta$ .

The circle pattern problem has a solution (unique up to scale) if and only if a coherent angle system exists [Rivin 1994; Bobenko and Springborn 2004]. A *coherent angle system* is an assignment of angles  $\hat{\alpha}_{ij}^k$  for all triangles such that

- (i) they are all positive

$$\hat{\alpha}_{ij}^k > 0,$$

- (ii) they sum to  $\pi$  in each triangle

$$\forall t_{ijk} \in T : \hat{\alpha}_{ij}^k + \hat{\alpha}_{jk}^i + \hat{\alpha}_{ki}^j = \pi,$$

- (iii) they satisfy Equations 1 (with  $\hat{\alpha}$  instead of  $\alpha$ ) for the given  $\theta$ -weights.

The solvability question for a circle pattern problem is therefore reduced to a linear feasibility problem with  $3|T|$  variables,  $3|T|$  inequality constraints and  $|T| + |E|$  equality constraints.

Conditions (i) and (ii) imply that all  $\hat{\alpha}_{ij}^k < \pi$ . Condition (iii) implies that the  $\hat{\alpha}_{ij}^k$  sum to  $2\pi$  around interior vertices and to  $\pi - \kappa_i$  for boundary vertices (with  $\kappa_i$  defined by Equation 4). This may give the false impression that finding a coherent angle system is equivalent to solving the circle pattern problem—in the sense that one could construct triangles with angles  $\hat{\alpha}_{ij}^k$  and lay them out. This is not so. The angles determine the triangles only up to scale. In general it is not possible to determine the size of each triangle in such a way that they all fit together. This observation is also what lead Sheffer and de Sturler [2000] to add their non-linear constraints on the product of sines of triangle angles.

Conditions (i), (ii) and (iii) combined imply that a necessary condition for the solvability of a circle pattern problem is that

$$\sum_{v_i \in V_{\text{bdy}}} \kappa_i = 2\pi, \quad (5)$$

with  $\kappa_i$  defined by Equation 4.

**Local Geometry of an Edge** To elucidate the role of terms which make up the variational energy characterizing the solution to the circle pattern problem we consider the local geometry around a given edge (see Figure 7). The basic building blocks of the parameterization are the *kites* formed by the endpoints of an edge ( $v_i, v_j$ ) and the incident face circumcircle centers  $c_{ijk}$  and  $c_{jil}$  for triangles  $t_{ijk}$  and  $t_{jil}$  respectively (at the boundary  $t_{jil}$  is missing). Since all relations are scale invariant it is convenient to introduce the *logarithmic* radius variables  $\rho_t = \log r_t$  for  $t \in T$ . With these definitions the angle  $\varphi_e^k$  induced at  $c_{ijk}$  by  $e_{ij}$  follows as

$$\varphi_e^k = \begin{cases} f_e(x) = \text{atan2}(\sin \theta_e, e^x - \cos \theta_e) & e \in E_{\text{int}} \\ \pi - \theta_e & e \in E_{\text{bdy}} \end{cases} \quad (6)$$

where  $x = \rho_{ijk} - \rho_{jil}$  (see Figure 7). The condition that every face in the parameterization should be flat constrains the  $\rho_t$  as functions of the  $\theta_e$  through the system of non-linear equations

$$\forall t \in T : 0 = 2\pi - \sum_{e \in t} 2\varphi_e^t, \quad (7)$$

i.e., summing around all edges incident on a face the resulting total angle must be  $2\pi$ .




Figure 7: Geometry around an edge (left). The kite formed by the edge endpoints and the incident face circumcircle centers ( $c_{ijk}, c_{jil}$ ) allows us to determine  $\varphi_e^k$  as a function of the given  $\theta_e$  and unknown radii  $r_{ijk}$  and  $r_{jil}$  (see Equation 6).

The basic idea of the variational energy formulation of Bobenko and Springborn now is to give an energy  $S(\rho)$  for which the face flatness Equations 7 are equivalent to the vanishing of the energy gradient,  $\nabla_\rho S = 0$ .

**Variational Characterization of Circle Patterns** To arrive at the desired energy Bobenko and Springborn used the fact that

$$F_e(x) = \int_{-\infty}^x f_e(\xi) d\xi = \text{Im Li}_2(e^{x+i\theta_e}),$$

where

$$\text{Li}_2(x) = - \int_0^x \frac{\log(1-\xi)}{\xi} d\xi = \sum_{k=1}^{\infty} \frac{x^k}{k^2} \quad |x| \leq 1$$

denotes the dilogarithm function. The imaginary part of the dilogarithm function of a complex argument can be expressed in terms of a  $2\pi$ -periodic real function (Clausen's integral) that can be computed efficiently with high accuracy. Using  $\rho_k$  and  $\rho_l$  as shorthand for  $\rho_{ijk}$  (respectively  $\rho_{jil}$ ) the energy is

$$S(\rho) = \sum_{e \in E_{\text{int}}} (\text{Im Li}_2(e^{\rho_k - \rho_l + i\theta_e}) + \text{Im Li}_2(e^{\rho_l - \rho_k + i\theta_e}) - (\pi - \theta_e)(\rho_k + \rho_l)) - \sum_{e \in E_{\text{bdy}}} 2(\pi - \theta_e)\rho_k + 2\pi \sum_{t \in T} \rho_t. \quad (8)$$

The gradient of this energy with respect to  $\rho_k$  is

$$\frac{\partial S}{\partial \rho_k} = 2\pi - \sum_{\{e \in k\} \cap E_{\text{int}}} 2f_e(\rho_k - \rho_l) - \sum_{\{e \in k\} \cap E_{\text{bdy}}} 2(\pi - \theta_e), \quad (9)$$

giving us the desired equivalence of  $\nabla_\rho S = 0$  and Equation 7. The Hessian of the energy is

$$d\rho^T (\text{Hess } S) d\rho = \sum_{e \in E_{\text{int}}} \frac{\sin \theta_e}{\cosh(d\rho_k - d\rho_l) - \cos \theta_e} (d\rho_k - d\rho_l)^2. \quad (10)$$

In particular from this expression we can see that the energy is convex except along the scaling “direction,” i.e., the Hessian has a null space spanned the constant vector  $d\rho = (1, 1, 1, \dots)$  and is otherwise positive. (This immediately implies the uniqueness of solutions of the circle pattern problem up to scale.)

## 2.2 Algorithm

There are three basic stages to the algorithm: (1) setting the  $\theta$  angles; (2) minimizing the energy; and (3) generating the layout. We discuss these in turn.



## 2.3 Edge Angles

As a first step of the algorithm,  $\theta_e$  angles need to be assigned to all edges of the mesh. These must satisfy the bounds constraints (Equation 2), sum conditions (Equation 3) at interior vertices, and, in the case of prescribed boundary curvatures, the boundary curvature conditions (Equations 4 and 5). Last but not least, a coherent angle system must exist for them. Of course, the  $\theta$  angles should also reflect the conformal structure of the original mesh as well as possible. Let  $\alpha_{ij}^k$  denote the angles in the original mesh. Ideally, one would like to assign  $\theta_e = \pi - \alpha_{ij}^k - \alpha_{ij}^l$ , where  $\alpha_{ij}^k$  and  $\alpha_{ij}^l$  are the angles opposite an interior edge  $e$  (and  $\theta_e = \pi - \alpha_{ij}^k$  for a boundary edge.) Then the intersection angles in the circle pattern would be the same as the intersection angles of the circumcircles in the mesh. This is too much to ask, though, because the conditions on  $\theta_e$  will be violated (after all the original mesh is not flat).

Our aim is to find angles  $\hat{\alpha}_{ij}^k$  close to  $\alpha_{ij}^k$  that can serve as a coherent angle system for the  $\theta$ -angles that we read off in the usual way. To this end, we minimize the objective function

$$\mathcal{Q}(\hat{\alpha}) = \sum |\hat{\alpha}_{ij}^k - \alpha_{ij}^k|^2$$

subject to the following constraints:

- positivity:  $\forall \hat{\alpha}_{ij}^k : \hat{\alpha}_{ij}^k > 0$ ,
- local Delaunay condition:  $\forall e_{ij} \in E_{\text{int}} : \hat{\alpha}_{ij}^k + \hat{\alpha}_{ij}^l < \pi$ ,
- triangle sum condition:  $\forall t_{ijk} \in T : \hat{\alpha}_{ij}^k + \hat{\alpha}_{jk}^i + \hat{\alpha}_{ki}^j = \pi$ ,
- vertex sum condition:  $\forall v_k \in V_{\text{int}} : \sum_{t_{ijk} \ni v_k} \hat{\alpha}_{ij}^k = 2\pi$ .

If we want natural boundary conditions (see for example the face in Figure 1 and the lion and Max Planck examples in Figures 9, 11), we add the constraint

- natural boundary condition:  $\forall v_k \in V_{\text{bdy}} : \sum_{t_{ijk} \ni v_k} \hat{\alpha}_{ij}^k < 2\pi$ .

If we want to prescribe the boundary curvature  $\kappa$  at boundary vertices (see the constrained boundary for the camel and the straight-line layout for Max Planck in Figure 1), we add the constraint

- prescribed boundary curvature:  $\forall v_k \in V_{\text{bdy}} : \sum_{t_{ijk} \ni v_k} \hat{\alpha}_{ij}^k = \pi - \kappa_k$ .

We solve this quadratic minimization problem with linear inequality constraints (on the  $\hat{\alpha}_{ij}^k$ ) with the software [Mosek 2005]. Since both the bounds constraints and the objective are convex and the additional constraints linear we have experienced no difficulty finding solutions efficiently even for large meshes (see Section 3).

Then we set  $\theta_e = \pi - \hat{\alpha}_{ij}^k - \hat{\alpha}_{ij}^l$  on interior edges and  $\theta_e = \pi - \hat{\alpha}_{ij}^k$  on boundary edges and proceed to the next stage of energy minimization.

It is theoretically possible that the constraints are not feasible. In the cases of natural boundary conditions and mapping to the disk this is due to the (counterintuitive) fact that there exist triangulations of the topological sphere that cannot be realized as convex polyhedra with vertices on the unit sphere; see, e.g., Grünbaum [2003]. These triangulations are rather special and we do not expect to encounter them in practice. In any case, it can be shown that 4 to 1 refinement applied once leads to a legal triangulation. In the case of prescribed boundary curvature, it may happen that the constraints become infeasible. We have not encountered any problems, but we have also not put any effort into exploring just how far one can go. (Our most complex example along these lines is the texture boundary for the

Camel with numerous prescribed curvatures all around the boundary.)

### 2.3.1 Mapping to the Disk

To map a mesh to the unit disk, we first pick a vertex at the boundary of the mesh and remove it together with all adjacent faces (see Figure 8). Now we map the resulting mesh to a convex polygon




Figure 8: Mapping a mesh to the disk. (1) Remove a boundary vertex together with all adjacent faces (dark pink). (2) Map the mesh to a convex polygon where all the original boundary edges lie on a straight line (top). The removed vertex (which should be imagined at infinity) and its incident triangles are shown schematically (dark pink). (3) Invert in a circle and reinsert the missing vertex to complete the mesh (bottom right). In this inversion we get to pick a vertex which will map to the center of the disk. In our example this vertex is the marked vertex between the ears of the cat.

in the plane with prescribed boundary curvature. At boundary vertices that were not adjacent to any of the removed faces, we fix the curvature to be 0. At the others, we bound the curvature to be positive. These conditions ensure that the original mesh boundary will be mapped to a straight line (Figure 8, top). Now we reflect in a circle whose center lies on the other side of that line. This maps the boundary line to a boundary circle. Finally we reinsert the removed vertex at the center of the reflection circle (it lies on the boundary circle) and complete the mesh (Figure 8, bottom right). Note that the circle reflection maps the other edges of the polygon to the circumcircles of the reinserted triangles. If we want to have a particular interior vertex at the center of the circular map, this can be achieved by an appropriate choice of reflection circle (or equivalently by a Möbius transformation).

In the polygonal mesh (Figure 8, top), the dynamic range of the edge lengths is typically large. This was initially a cause for concern, but it turned out to be harmless. We give a rough argument why this should be so. Suppose the mesh has  $n$  vertices. Then the number of boundary edges will be roughly  $n^{\frac{1}{2}}$ . Suppose that in the final circular map the boundary vertices are evenly spaced on the boundary circle. After inversion on a circle that sends one of these vertices to infinity, one finds that the largest (finite) edge on the

boundary line is proportional to  $n^{\frac{1}{2}}$  while the smallest edge is proportional to  $n^{-\frac{1}{2}}$ . Hence the ratio is  $n$ . If  $n = 10^6$  we would expect to lose about 6 out of 16 double precision digits. Figure 11 shows an example of a larger mesh mapped to the disk. We experienced no numerical difficulties in this procedure.

**Discussion** Angle optimization is also at the heart of the work of Sheffer and de Sturler [2000], who formulate their flattening problem as one which minimizes the (weighted) least squares deviation of the measured  $\alpha_{ij}^k$  (scaled to satisfy the angle sum condition around a vertex) from the realizable  $\hat{\alpha}_{ij}^k$  with flatness sum conditions for each triangle and each vertex. Unfortunately, when using the  $\hat{\alpha}_{ij}^k$  to determine the final shape of the triangles in the flat mesh, one must additionally include non-linear conditions on the quotients of sines of  $\hat{\alpha}_{ij}^k$  around each vertex (due to the law of sines). This additional non-linear condition makes the minimization problem numerically much harder since it becomes non-convex (for recent significant progress in the numerical treatment see [Sheffer et al. 2004]). Due to the lack of convexity, it cannot be expected that there is only one local minimum.

Our method works in two stages. First we change the angles of the triangles, but only so much that we can read off valid  $\theta$ -angles. The deviation from the ideal angles is measured with a quadratic objective just as in the work of Sheffer and de Sturler. The critical observation is that the sum of two  $\alpha$ -angles across an edge is a conformal invariant (because it measures the angle of intersection between the circumcircles), while the  $\alpha$ -angles themselves are not. In other words, the best one can hope for in a discrete conformal mapping is conservation of  $\theta$ -angles, not  $\alpha$ -angles. Formulating the problem in terms of given  $\theta$ -angles then leads to a simple convex, and thus *unique*, minimization problem: the final circle pattern in the plane is completely determined by the  $\theta$ -angles.

An alternative approach for the construction of mappings to the disk in the context of discrete harmonic mappings (cotan formula) was proposed by Jin and co-workers [2004]. They “welded” a mesh with its double at the boundary to create a sphere topology and then proceeded to map to the sphere. Enforcing symmetry in this mapping turns the original boundary into an equator. They had to effectively double the degrees of freedom while in our approach the number of degrees of freedom is (essentially) constant. As they observed, being able to map to a canonical region such as a disk one can effectively map from any region to any other through composition of one map to the disk with the inverse of another map to the disk (though the meshes in general do not match up and some interpolation in the disk is required).

## 2.4 Energy Minimization

With a given valid assignment of  $\theta_e$  for all edges the energy minimization is straight forward (using Equations 8, 9, and 10). We have relied on a black box energy minimizer [Mosek 2005] with excellent results (see Section 3).

## 2.5 Layout Generation

Once all radii have been determined the length of each edge follows easily from a *local* computation

$$|e_{ij}| = 2r_k \sin \varphi_e^k = 2r_l \sin \varphi_e^l$$




Figure 9: Comparison of two different mappings of the cut Max Planck model: natural boundaries and a straightline boundary. Error plots indicate the quasiconformal distortion between original and mapped triangles.

(see Figure 7). We begin the layout with some interior edge, starting it at the origin and orienting it along the positive  $x$ -axis and push it onto an empty stack. When popping an edge off the stack we lay out any vertices not yet laid out in the incident triangles (at most two). This results in at most four edges being finished (second end point fixed); such edges are pushed onto the stack. The process repeats until the stack is empty. There may be concern that such a procedure might lead to cumulating error as one proceeds. We have found consistently that we achieve accuracies on the order of  $10^{-8}$  with double precision data (for example for the Camel, Max Planck, and Lion datasets). Should accuracy of the layout become an issue we recommend the procedure given in [Sheffer et al. 2004], which solves for all vertex positions simultaneously with a simple (least squares induced) linear system.

## 3 Results

Figure 1 demonstrates different boundary shapes that can be obtained by our method. The Max Planck head is parameterized by a simple polygonal region. This was achieved by prescribing the boundary curvature of the parameter domain. It was set to zero at all boundary vertices, except for eight designated corner vertices, where it was set to  $\pm\pi/2$ . The parameterized camel shows that more complicated boundary shapes can also be achieved by prescribing the boundary curvature. Here the parameter domain is essentially a polygonal region with rounded corners. Both the Max



Figure 10: Our method is robust to varying sampling rates. Here a symmetric (left/right) geometry sampled at different rates. The parameterization (with natural boundary) maintains the symmetry of the geometry.

Planck mesh and the camel mesh were cut before parameterization. We do nothing to ensure continuity across the cut and, unsurprisingly, there is none. The face was parameterized with natural boundary conditions and the lion head was mapped to a disk. The following table shows timings for the angle optimization and energy minimization. The layout times are negligible ( $\sim 0.5$ s). The timings were measured on a 3.6GHz Pentium IV Xeon.

Model	Faces	Angles	Energy
Max Planck	37.9K	26s	9s
Camel	77.7K	56s	17s
Face	12.6K	7s	2s
Lion Head	39.6K	28s	8s

We claim that our discrete conformal maps are close to angle preserving. A measure for the conformality of a map is the *quasiconformal distortion*: the ratio of the larger to the smaller eigen value of the Jacobian matrix. It is at least 1 and equal to 1 everywhere only for conformal maps. Figures 9 and 11 show the quasiconformal distortion of our discrete conformal maps for different meshes and different boundary conditions. For the Max Planck head with natural boundary conditions, the average and maximal values are 1.02 and 1.35. With the polygonal boundary conditions the maximum distortion goes up to 4.06, but the high distortion is concentrated at the boundary and spreads very little inwards. Even the example of the highly convoluted lion head (Figure 11) shows very low conformal distortion in most places with a few hot spots of high distortion. Because the natural boundary shape is already fairly round, the difference between natural boundary conditions and disk boundary is not so pronounced.

Figure 10 shows the results of an experiment to see how sensitive our method is to abruptly differing sampling rates. The geometry of the head is symmetric while the sampling rate doubles at the right/left symmetry line. Examining the flattened mesh (using natural boundary conditions) we observe that the left/right symmetry is preserved (see also the resulting texture mapping on the original surface).




Figure 11: Comparison of two different mappings of the Lion data set: natural boundaries and disk boundary. Error plots indicate the quasiconformal distortion between original and mapped triangles.

## 4 Conclusion

We have presented a new method to parameterize surface meshes. It is based on the mathematical theory of *circle patterns* and produces discrete conformal maps. The shape of the boundary may be determined by natural boundary conditions or by prescribing the curvature of the boundary. This affords a high degree of control over the boundary shape ranging from disks and simple polygonal outlines to more complex boundary arrangements. The examples show that our method is efficient, provides good parameterizations even for large and complex meshes, and that the result is insensitive to the way the surface is triangulated.

Our algorithm works in three stages: first, we solve a quadratic programming problem to obtain intersection angles. These are the input for the second stage that consists in the minimization of a convex energy. The output of this stage (the radii of the circumcircles) and the intersection angles determine the shape of all triangles which are then laid out. In every stage, the solution is unique and depends continuously on the input.

At the moment, we use general purpose solvers in the first two stages, and a very simple layout algorithm in the third. Efficiency could be improved by customizing and tuning the minimizers, and by switching to hierarchical methods [Sheffer et al. 2004]. Their sophisticated layout scheme—it minimizes the global layout error—can also be used without change as the third stage in our algorithm.

When we parameterize a mesh with prescribed boundary curvature, we have to set the curvature at each boundary vertex. At present, we do this manually. A user friendlier graphical interface is desirable. Since it is most likely that the performance of our method can be improved further, it is not unreasonable to envision an interactive interface that lets the user manipulate the parameter domain while the parameterization is incrementally recomputed.

The most exciting avenue for future work concerns the use of variational methods for the construction of discrete conformal mappings to the sphere and for surfaces of higher genus. Energy functionals for circle patterns on the sphere as well as hyperbolic space exist and can be used in a manner similar to our approach presented here. There are both additional difficulties (the sphere functional is *not* convex) and opportunities: the hyperbolic circle pattern case does *not* require a surface with higher genus to be cut open first. Instead one can solve the energy minimization directly and only after the solution has been found (through a layout in the hyperbolic plane) pick a suitable fundamental domain.

**Acknowledgments** This work was supported in part by NSF (DMS-0220905, DMS-0138458, ACI-0219979), DFG (Research Center MATHEON “Mathematics for Key Technologies,” Berlin), DOE (W-7405-ENG-48/B341492), nVidia, the Center for Integrated Multiscale Modeling and Simulation, Alias, and Pixar. Special thanks to Alexander Bobenko, Mathieu Desbrun, Ilja Friedel, and Cici Koenig.

## References

- BOBENKO, A. I., AND SPRINGBORN, B. A. 2004. [Variational Principles for Circle Patterns and Koebe’s Theorem](#). *Transactions of the American Mathematical Society* 356, 659–689.
- BOWERS, P. L., AND HURDAL, M. K. 2003. [Planar Conformal Mappings of Piecewise Flat Surfaces](#). In *Visualization and Mathematics III*, Springer-Verlag, Berlin, H.-C. Hege and K. Polthier,



- Eds., Mathematics and Visualization, 3–34. Papers from the 3rd International Workshop held in Berlin, May 22–25, 2002.
- BRÄGGER, W. 1992. Kreispackungen und Triangulierungen. *Enseign. Math.* 38, 201–217.
- COLIN DE VERDIÈRE, Y. 1991. Un principe variationnel pour les empilements de cercles. *Invent. Math.* 104, 655–669.
- COLLINS, C., AND STEPHENSON, K. 2003. [A Circle Packing Algorithm](#). *Computational Geometry: Theory and Applications* 25, 233–256.
- DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. [Intrinsic Parameterizations of Surface Meshes](#). *Computer Graphics Forum (Proceedings of Eurographics 2002)* 21, 3, 209–218.
- GRÜNBAUM, B. 2003. *Convex polytopes*, second ed., vol. 221 of *Graduate Texts in Mathematics*. Springer-Verlag, New York.
- GU, X., AND YAU, S.-T. 2003. [Global Conformal Surface Parameterization](#). In *Symposium on Geometry Processing*, Eurographics/ACM SIGGRAPH, 127–137.
- HIRANI, A. N. 2003. [Discrete Exterior Calculus](#). PhD thesis, California Institute of Technology.
- JIN, M., WANG, Y., YAU, S.-T., AND GU, X. 2004. [Optimal Global Conformal Surface Parameterizations](#). In *Proceedings of IEEE Visualization*, IEEE, 267–274.
- LEIBON, G. 2002. Characterizing the Delaunay Decompositions of Compact Hyperbolic Surfaces. *Geom. Topol.* 6, 361–391.
- LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. [Least Squares Conformal Maps for Automatic Texture Atlas Generation](#). *ACM Transactions on Graphics* 21, 3, 362–371.
- MERCAT, C. 2001. [Discrete Riemann Surfaces and the Ising Model](#). *Communications in Mathematical Physics* 218, 1, 177–216.
- MOSEK, 2005. Constrained Quadratic Minimization Software. <http://www.mosek.com/>. Version 3.1r42.
- PINKALL, U., AND POLTHIER, K. 1993. [Computing Discrete Minimal Surfaces and Their Conjugates](#). *Experimental Mathematics* 2, 1, 15–36.
- RIVIN, I. 1994. [Euclidean Structures on Simplicial Surfaces and Hyperbolic Volume](#). *Annals of Mathematics* 139, 3, 553–580.
- RODIN, B., AND SULLIVAN, D. 1987. The Convergence of Circle Packings to the Riemann Mapping. *J. Differential Geom.* 26, 2, 349–360.
- SHEFFER, A., AND DE STURLER, E. 2000. [Surface Parameterization for Meshing by Triangulation Flattening](#). In *Proceedings of the 9th International Meshing Roundtable*, Sandia National Labs, 161–172.
- SHEFFER, A., LÉVY, B., MOGILNITSKI, M., AND BOGOMYAKOV, A. 2004. [ABF++: Fast and Robust Angle Based Flattening](#). *ACM Transactions on Graphics*. To appear.
- THURSTON, W. P. 1980. [The Geometry and Topology of Three-Manifolds](#). Available at <http://www.msri.org/publications/books/gt3m/>.
- THURSTON, W. P. 1985. The finite Riemann mapping theorem. Invited talk at the symposium on the occasion of the proof of the Bieberbach conjecture held at Purdue University, March 1985.