

Manifold splines with a single extraordinary point

Xianfeng Gu^a, Ying He^{b,*}, Miao Jin^a, Feng Luo^c, Hong Qin^a, Shing-Tung Yau^d

^a Computer Science Department, Stony Brook University, NY, USA

^b School of Computer Engineering, Nanyang Technological University, Singapore

^c Mathematics Department, Rutgers University, NJ, USA

^d Mathematics Department, Harvard University, MA, USA

Received 19 August 2007; accepted 26 January 2008

Abstract

This paper develops a novel computational technique to define and construct manifold splines with only one singular point by employing the rigorous mathematical theory of Ricci flow. The central idea and new computational paradigm of manifold splines are to systematically extend the algorithmic pipeline of spline surface construction from any planar domain to an arbitrary topology. As a result, manifold splines can unify planar spline representations as their special cases. Despite its earlier success, the existing manifold spline framework is plagued by the topology-dependent, large number of singular points (i.e., $|2g - 2|$ for any genus- g surface), where the analysis of surface behaviors such as continuity remains extremely difficult. The unique theoretical contribution of this paper is that we devise new mathematical tools so that manifold splines can now be constructed with only one singular point, reaching their theoretic lower bound of singularity for real-world applications. Our new algorithm is founded upon the concept of discrete Ricci flow and associated techniques. First, Ricci flow is employed to compute a special metric of any manifold domain (serving as a parametric domain for manifold splines), such that the metric becomes flat everywhere except at one point. Then, the metric naturally induces an affine atlas covering the entire manifold except this singular point. Finally, manifold splines are defined over this affine atlas. The Ricci flow method is theoretically sound, and practically simple and efficient. We conduct various shape experiments and our new theoretical and algorithmic results alleviate the modeling difficulty of manifold splines, and hence, promote the widespread use of manifold splines in surface and solid modeling, geometric design, and reverse engineering.

© 2008 Elsevier Ltd. All rights reserved.

Keywords: Manifold splines; Affine structure; Discrete Ricci flow; Extraordinary point; Metric; Differential geometry

1. Introduction and motivation

1.1. Problem statement

Despite many algorithmic and theoretical advances in solid modeling and shape computing in most recent years, one fundamental objective of our research community has always been to develop novel modeling, design, and simulation schemes that are capable of accurately representing complicated real-world objects in a compact manner, and facilitating rapid computation of their desirable properties both

globally and locally such as differential properties, smoothness requirements, and topological validity. Strongly inspired by the recent development of subdivision surfaces and manifold splines, our current research goal in this paper is to further advance the state of the knowledge about manifold splines. At the theoretic level, we devise manifold splines with only one singular point through the mathematical rigor of Ricci flow and relevant computational techniques. At the application level, we design a brand new algorithmic pipeline that enables all the computational elements needed for the widespread use of manifold splines (especially the new, improved scheme with a single extraordinary point) in solid modeling, shape design, and reverse engineering.

1.2. Manifold splines

For the perspectives of solid modeling, engineering design, finite element simulation, and scientific computation, elegant

* Corresponding address: 50 Nanyang Avenue, BLK N4, Nanyang Technological University, 639798, Singapore. Tel.: +65 65141008; fax: +65 67926559.

E-mail addresses: gu@cs.sunysb.edu (X. Gu), yhe@ntu.edu.sg (Y. He), mjin@cs.sunysb.edu (M. Jin), fengluo@math.rutgers.edu (F. Luo), qin@cs.sunysb.edu (H. Qin), yau@math.harvard.edu (S.-T. Yau).

geometric properties such as high-order continuity and the ease of computing all the desirable properties rapidly are always mandatory for the development of novel shape representations. Therefore, it is not surprising to see that spline-centric polar forms [21] are becoming the most popular computational tools in geometric modeling and shape design. Essentially, the methodology of polar forms naturally gives rise to parameterization-centered, piecewise polynomials defined on any planar parameter domain for the effective modeling and accurate computing of smooth spline surfaces.

However, examining all the real-world applications, we observe that the most natural shapes are manifolds with complicated topologies and arbitrarily detailed geometric configurations, which cannot be completely covered by a single open surface defined in one coordinate system (note that, it does not matter if the parametric surface is a polynomial or a non-polynomial; this fundamental principle remains the same). Instead, a manifold might be covered by a family of *coordinate charts*, each coordinate chart covering only a portion of the manifold. Different charts may overlap with each other, a *coordinate transition function* transforms from one coordinate system to the other. If we follow the algorithmic procedure of polar forms and other relevant computational techniques in a principled way, we can easily realize that conventional splines (defined over any open domain) cannot be transferred over the manifold directly.

In order to model a manifold using piecewise polynomials, current approaches will segment the manifold into many patches, define a single coordinate system over each patch, such that each patch can be modeled by a spline patch. Finally, any generic approach will glue/about all the spline patches together by adjusting the control points and the knots along their common boundaries. This whole process is mainly performed manually, and it requires the users' skill and mathematical sophistication, and is tedious and error-prone.

It is highly desirable to design splines defined over manifolds directly, such that different spline patches can be automatically glued together with high continuity, so the modelling process requires neither segmentation nor patching. Pioneering work has been done by Grimm and Hughes [6], which can model splines on arbitrary surfaces. Recently, Ying and Zorin [25] introduced a general method by constructing a conformal atlas. In both methods, smooth functions are defined on each chart and blended together to form a function coherently defined over the entire manifold. The methods are flexible for all manifolds with arbitrary topologies. The functions are created with any degree of desirable continuity without any singularity. The primary drawbacks of these methods are that surfaces constructed this way are no-longer polynomials and their computational expenses are relatively high in comparison with conventional spline surfaces.

Most recently, the manifold splines proposed by Gu, He, and Qin [8] offer a different approach to manifold domain construction. The main advantage for manifold splines is that on each local chart, the functions are all piecewise polynomials in common use; currently available spline surfaces in commercial software packages can easily serve as building

blocks towards the effective design of complicated models of arbitrary topology, and hence, the evaluation and all the necessary computational procedures are both efficient and robust. Furthermore, existing algorithms and software for conventional spline surfaces can all be easily adopted for use in applications of manifold surfaces. Nevertheless, certain drawbacks still remain: there must be singularities for general manifolds except tori. In [8], they discovered that the existence of manifold splines is equivalent to the existence of a special atlas of the underlying manifold domain, whose transition functions are all affine among themselves, denoted as their *affine atlas*. Unfortunately, it is impossible to find an affine atlas to cover arbitrary closed surfaces except tori. There must be singularities for the atlas which cannot be covered by any chart within its collection set. Moreover, they proved that the minimal number of singularities equals one without developing any practical algorithm. So, how to lower the number of singular points remains elusive, and how to devise new algorithms with a minimum number of singular points for practical applications remains extremely technically challenging.

Using existing popular techniques of polar forms and Euclidean metrics, the manifold splines in the neighborhood of singularities appear to be extremely difficult to construct, unstable, and error-prone. In addition, the mapping distortion from the surface to the affine atlas significantly affects the quality of the final spline surface. The distortions are intrinsically determined by the singularities for the affine atlas. Therefore, it is highly desirable for users to be able to control the positions and the number of singularities. For open surface cases, however, it is ideal to push the singularities away from the surfaces. In this paper, we demonstrate that the Discrete Ricci flow is a powerful theoretic and computational tool for constructing an affine atlas with full control of singularities, and specifically, one capable of minimizing the number of singularities to its theoretical lower bound (which is at most one for closed surfaces and zero for all open surfaces).

1.3. Intrinsic shape space

In reality, surfaces are typically acquired via modern scanning devices, and they are initially approximated by a set of points and/or triangular meshes. We shall consider the triangle-mesh approximation of domain manifolds first. In order to find an affine atlas of a triangle mesh, it is sufficient to find a configuration of edge lengths such that the one-ring neighbor of each vertex is flat. So, any parameterization problem can be formulated as:

Finding a configuration of edge lengths, such that each vertex has zero discrete Gaussian curvature¹

One could naturally raise the following much broader questions: given a mesh,

- (i) *What are the all possible configurations of edge lengths?*

¹ Discrete Gaussian curvature is defined as the difference between 2π and the summation of all angles adjacent to the vertex (see Section 3.2).

- (ii) What are the all possible configurations of curvatures on vertices?
- (iii) What is the relation between edge length configurations and curvature configurations? It is obvious that edge lengths determine curvatures. Can curvatures determine edge lengths?

The entire space of all possible configurations of edge lengths is denoted as the *metric space*. The entire space of all possible configurations of vertex curvatures is denoted as the *curvature space*. Metric space and curvature space are *intrinsic shape spaces* of the mesh.

The answer to the admissible edge lengths is straightforward: any configuration satisfying triangle inequality is admissible. For admissible curvature configuration, the answer is much more complicated. There are mainly two constraints: topological constraint and combinatorial constraint. The topological constraint is represented as the Gauss–Bonnet formula, and the total curvature equals to the $2\pi\chi$, where χ is the Euler number of the mesh. The combinatorial constraint ensures that all angles are between 0 and π , and represented solely by inequalities of curvature and connectivity. The technical details will be discussed in the next Section.

The answer to the third question has fundamental importance, it is the main focus of this research work. It is easy to compute curvature using edge lengths, but the inverse is much more complicated. Intuitively speaking, the Gaussian curvature is a map from the metric space to the curvature space, and the mapping in general is not injective. However, one can select a subspace of the whole metric space, such that any two metrics in the subspace are conformally equivalent. Restricted to this subspace, Gauss curvature map is a homeomorphism, namely, any curvature configuration uniquely determines an edge length configuration in this subspace.

In practice, one can specify the target curvature and deform the edge length according to the difference between the current curvature and the target curvature. It is guaranteed that the curvature configuration of the final mesh will reach the target one. This kind of deformation process driven by Gaussian curvature is the so called *Ricci flow*.

1.4. Ricci flow

Ricci flow was first introduced in differential geometry by Hamilton in [9]. It has solid theoretical foundations. By nature, it is a constructive geometric tool and can be easily implemented; therefore, it has a great potential for real-world applications.

The fundamental idea of Ricci flow is rather simple. We can deform the surface driven by its curvature to the desired shape. Suppose S is a closed surface with Riemannian metric \mathbf{g} , and u is a function on S ; then $e^{2u}\mathbf{g}$ is another metric on S conformal to \mathbf{g} . Ricci flow is explicitly defined as

$$\frac{du(t)}{dt} = \bar{K} - K(t), \quad (1)$$

where the area preserving constraint is explicitly formulated as

$$\int_S dA = \int_S e^{2u} dA, \quad (2)$$

and $K(t)$ is the Gaussian curvature induced by the metric $e^{2u(t)}\mathbf{g}$, and \bar{K} is a constant

$$\bar{K} = \frac{2\pi\chi(S)}{\int_S dA}.$$

It has been proven that Ricci flow converges to the uniform metric that induces constant Gaussian curvature \bar{K} on the surface, $K(\infty) \rightarrow \bar{K}$. Furthermore, Ricci flow converges to the final stable solution exponentially fast: for a given surface S , there exist two positive constants c_1, c_2 determined by the geometry of S , such that,

$$|K(t) - K(\infty)| < c_1 e^{-c_2 t}.$$

Ricci flow has many promising properties, which make it very valuable for real-world applications,

- Ricci flow offers the freedom to traverse the intrinsic shape space (all the admissible configurations of edge lengths) by driving the surface to deform to all possible shapes as long as the Gaussian curvature of the target shape is known.
- The deformation induced by Ricci flow is conformal.
- Ricci flow deforms the surface to a single solution and converges to the solution exponentially fast.
- Ricci flow can be formulated as a variational problem. The energy is a convex function; therefore it has a single global optimum. Ricci flow is the negative gradient flow of the energy, and can be further speed up using Newton's method.

In our current research, the fundamental motivation for us to use Ricci flow is its computational power to compute the affine atlas of a mesh with any desired number of singularities, especially with only one singularity.

1.5. Contributions

In this paper, we devise a novel algorithm to construct manifold splines with only one singular point for closed surfaces and no singular point for open surfaces, reaching its lower bound in theory. The algorithm is uniquely founded up on a mathematically rigorous tool in differential geometry, namely, Ricci flow. Key contributions of this paper include:

- (i) We formulate the intrinsic space of a mesh: the metric space (i.e., all admissible configurations of edge lengths) and the curvature space (i.e., all admissible configurations of vertex curvatures). We point out the topological constraints and the combinatorial constraints for the metric spaces. We re-define the general surface parameterization problem as being equivalent to finding flat metrics with any user-assigned singularities.
- (ii) We articulate our new computational method to construct an affine atlas with any pre-determined singularities using Ricci flow. The affine atlases serve as the key and necessary elements for constructing manifold splines, especially, for manifold splines with only one singular point.

- (iii) We offer a theoretically rigorous, practically simple and computationally efficient tool, Ricci flow, to solve geometric and solid modeling problems. In its discrete setting, given the Gaussian curvature on each vertex, Ricci flow will be employed to compute the configuration of edge lengths.

2. Background review

This section briefly reviews previous work on splines, Ricci flow, parameterization, and circle packing.

Manifold splines. Pioneering work has been done earlier by Grimm and Hughes [6], which can model splines on arbitrary surfaces. Recently, Ying and Zorin [25] introduced a general method by constructing a conformal atlas. The function bases in their constructions are smooth and without singularities; however, they are not polynomials, requiring the necessary data exchange between polynomial-based spline surfaces and their special-purpose functions for surface design.

Manifold splines defined by piecewise polynomials over manifold domains of arbitrary topology were first rigorously formulated in [8], which unifies the conventional spline surfaces based on polar forms and the subdivision surfaces of arbitrary topology. In their work, it is proven that a manifold admits manifold splines based on polar forms if and only if it has an affine atlas. The topological obstruction for the existence of the affine atlas is the Euler class. By removing only one point, any oriented surface has an affine atlas.

Ricci flow. Ricci flow on a surface is introduced by Hamilton in [9], which will conformally deform the metric of a surface to a canonical metric with constant Gaussian curvature. For a closed genus-zero surface, Ricci flow will change the metric to the spherical metric with constant positive Gaussian curvature; for a genus-one closed surface, the solution to Ricci flow is the planar metric with zero Gaussian curvature; for a high genus closed surface, the solution to Ricci flow is the hyperbolic metric with constant negative Gaussian curvature. The analogue of Ricci flow in the discrete, combinatorial setting is first studied in [3]. It is proven that combinatorial Ricci flow will deform the metric of a triangle mesh to metrics with constant vertex curvatures. Recently, Jin et al. applied discrete Ricci flow to compute the hyperbolic and real projective structures of surfaces [16].

Circle packing and circle pattern. Circle packing and circle patterns are used for approximating conformal deformations. Circle packing is first introduced by Thurston in [24], where he designed an algorithm to find the circle packing of a graph by adjusting the radii at vertices one at a time. Stephenson et al. developed practical algorithms in [23]. Circle pattern is introduced in [1] and applied for surface parameterizations in [17], which is closely related to circle packing. Instead of using circles centered at each vertex, this method uses the circumcircles of triangles. Comparing with circle pattern, the theoretic framework of Ricci flow is much simpler and clearer. Furthermore, the implementation of the Ricci flow is much easier in practice.

Global surface parametrization. The affine atlas can be computed using surface parametrization algorithms. In the literature, there exist many parameterization methods using a variety of distortion metrics. For a thorough survey, we refer the readers to the excellent work of Floater and Hormann [5, 22]. We shall focus on the most related work, especially global parameterization methods.

Gu and Yau computed the conformal structure based on Hodges theory in [7]. The method computes the holomorphic 1-form basis, and induces a flat metric with $2g - 2$ singularities. Ni et al. extracted the topological structure using a harmonic Morse function; the vector fields are holomorphic 1-forms, and induce a flat metric with more singularities [18]. Recently, Ray et al. [19] computed the global conformal parameterization also using holomorphic 1-forms, but specifically tailored the parameterization to follow the principle curvature lines.

It may be noted that, all current parameterization methods will introduce **multiple** singularities due to the topological obstruction. **The method to be developed in this paper is capable of reducing the number of singularity points to its theoretical lower bound (which is one).**

3. Global surface parameterization using discrete Ricci flow

Conventional *local* surface parameterization refers to the process of mapping a simply connected surface patch to a planar region. In contrast, a *global* surface parameterization maps the whole surface to the plane \mathbb{R}^2 , the unit sphere \mathbb{S}^2 or the hyperbolic space \mathbb{H}^2 periodically. The global surface parameterization problem could be formulated in a precise and general way as deforming the given surface to satisfy the prescribed curvatures. By deforming the surface, we mean finding a different Riemannian metric (the first fundamental form). If conformality is required, then the new metric should be conformal to the original metric. Mathematically, suppose \mathbf{g} is the original metric; then the metric conformal to \mathbf{g} has the form $e^{2u}\mathbf{g}$, where u is the function defined on the surface. Then global surface parameterization is to solve function u by the prescribed curvature.

In the following, we assume the surface is an oriented 2-manifold, represented by a two dimensional simplicial complex (i.e., triangular mesh) $M = (V, E, F)$, where V is the set of all vertices, E is the set of all non-oriented edges, and F the set of all faces. We use $v_i, i = 1, 2, \dots, n$ to denote its vertices, e_{ij} to denote an oriented edge from v_i to v_j , f_{ijk} to denote an oriented face, v_i, v_j, v_k are sorted counterclockwise.

3.1. Discrete conformal metrics

The central task is to approximate Ricci flow (1) in the discrete mesh setting. Continuous Ricci flow conformally deforms a surface.

Fig. 1 illustrates an important observation for continuous conformal mappings: they transform infinitesimal circles to infinitesimal circles, and preserve the intersection angles among the circles. Based on this property, Thurston introduced the

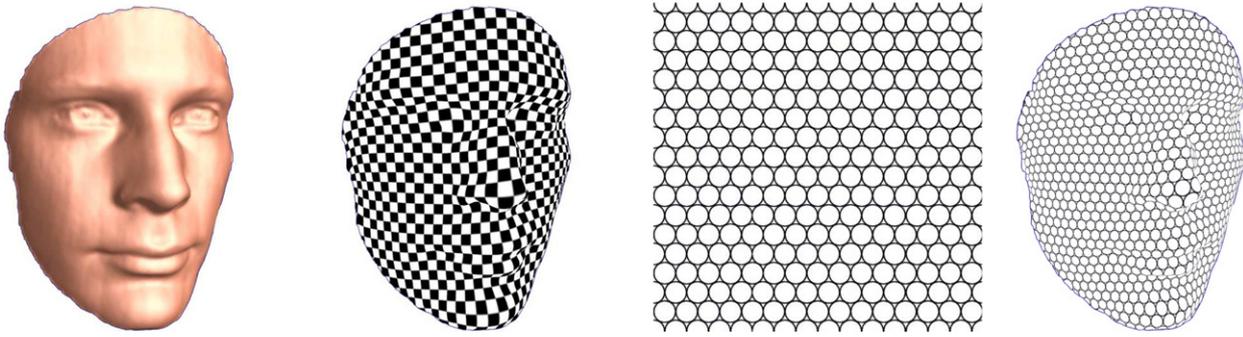
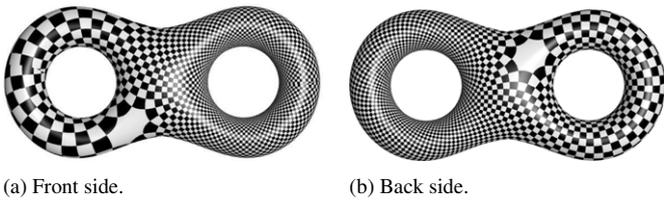


Fig. 1. Circle packing for a surface. Conformal mappings transform the infinitesimal circles on the texture plane to the infinitesimal circles on the surface.



circle packing metric in early eighties [24]: a circle with the radius γ_i is associated with each vertex v_i . For an each edge e_{ij} , two circles intersect at the angle Φ_{ij} , called the edge weight. The edge length of e_{ij} is determined by γ_i, γ_j and Φ_{ij} ,

$$l_{ij} = \sqrt{\gamma_i^2 + \gamma_j^2 + 2\gamma_i\gamma_j \cos \Phi_{ij}}. \tag{3}$$

Fig. 2. Affine atlas induced by a global conformal surface parameterization. The affine atlas is illustrated by the texture mapping of a checkerboard pattern. There are $2g - 2$ singularities centered at the white octagons.

It can be shown that for any face f_{ijk} with vertex radii $\{\gamma_i, \gamma_j, \gamma_k\}$ and edge weights $\{\Phi_{ij}, \Phi_{jk}, \Phi_{ki}\}$, if edge weights are acute angles, then the edge lengths $\{l_{ij}, l_{jk}, l_{ki}\}$ satisfy the

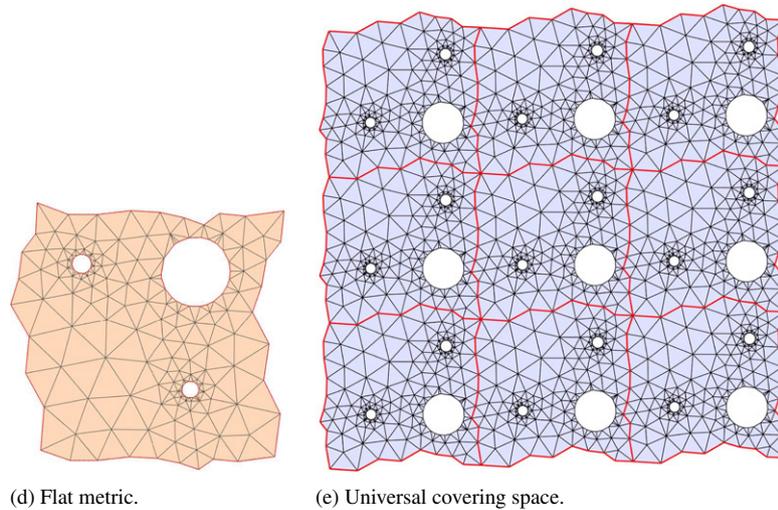
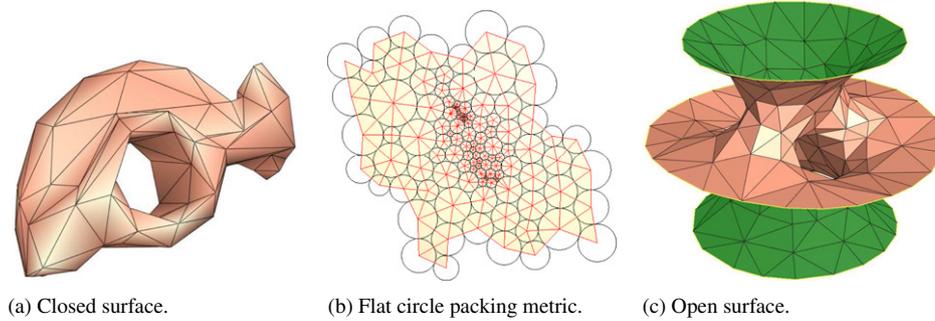


Fig. 3. Computing the affine structures for genus one surfaces using discrete Ricci flow. The right column (b) and (d) shows the embedded fundamental domain. The last row shows the universal covering space.

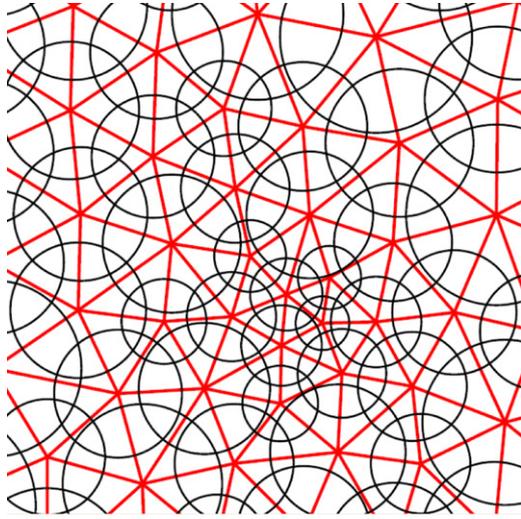


Fig. 4. Close-up view of the flat circle packing metric using Ricci flow.

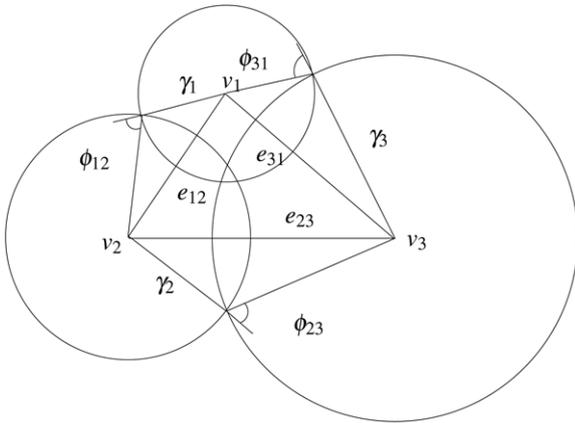


Fig. 5. Circle packing metric for a triangle. Triangle $[v_1, v_2, v_3]$ has vertices v_1, v_2 and v_3 , edges e_{12}, e_{23} and e_{31} . Three circles centered at v_1, v_2, v_3 , with radii γ_1, γ_2 and γ_3 intersect one another, the intersection angles are ϕ_{12}, ϕ_{23} and ϕ_{31} , which are the weights associated with the edges. The edge lengths of the triangle are determined by γ_i and ϕ_{ij} by the cosine law.

triangle inequality,

$$l_{ij} + l_{jk} > l_{ki}.$$

We use $\Gamma : V \rightarrow \mathbb{R}^+$ to denote the vertex radii, $\Phi : E \rightarrow [0, \frac{\pi}{2}]$ for the edge weight, then a circle packing metric is represented as (M, Γ, Φ) .

Two circle packing metrics (M, Γ_1, Φ_1) and (M, Γ_2, Φ_2) are conformal to each other, if $\Phi_1 \equiv \Phi_2$. Namely, a discrete conformal mapping will change the vertex radii only and preserve the intersection angles. Figs. 4 and 5 illustrate the circle packing metric.

3.2. Discrete curvature

Given a discrete metric (M, Φ, Γ) , suppose f_{ijk} is a face, the angle of vertex v_i in f_{ijk} is denoted as θ_i^{jk} ; then the discrete Gaussian curvature K_i at an interior vertex v_i is defined as

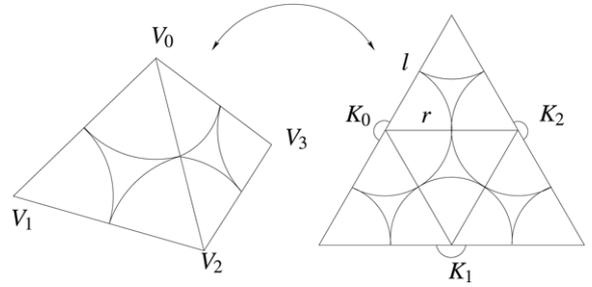


Fig. 6. Circle packing metric and curvature. For a canonical tetrahedron, the edge lengths are equal to $l = 1.0$, the radii on all the vertices are equal to $r = 0.5$. The curvature on each vertex equals to $K_i = \pi$. The weights of all edges Φ equal 0.

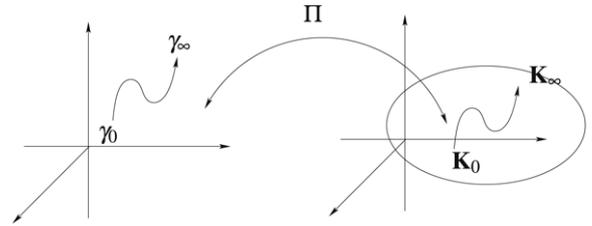


Fig. 7. Gaussian curvature is a homeomorphism between the circle packing metric space based on (M, Φ) and the curvature space, the inverse map can be computed using Ricci flow. We start from the known metric γ_0 and the known curvature \mathbf{k}_0 , then flow to the target curvature \mathbf{K}_∞ using Ricci flow, then the metric will flow to the corresponding metric $\gamma_\infty = \Pi^{-1}(\mathbf{K}_\infty)$.

$$K_i = 2\pi - \sum_{f_{ijk} \in F} \theta_i^{jk}, v_i \notin \partial M, \tag{4}$$

The discrete Gaussian curvature for an boundary vertex v_i is defined as

$$K_i = \pi - \sum_{f_{ijk} \in F} \theta_i^{jk}, v_i \in \partial M. \tag{5}$$

Fig. 6 demonstrates the circle packing metric for a tetrahedron surface, where all the edge weights are zeros, all the vertex radii are 0.5, and all the vertex curvatures are π .

The Gaussian curvature at each vertex could be arbitrary, but the total curvature is confined by the topology of the surface. This is indicated by the Gauss–Bonnet theorem.

Theorem 1 (Gauss–Bonnet). Suppose M is a mesh; the total discrete Gaussian curvature equals the product of 2π and its Euler number,

$$\sum K_i = 2\pi \chi. \tag{6}$$

Furthermore, for any discrete metric (M, Φ, Γ) , $\Phi : E \rightarrow [0, \frac{\pi}{2}]$ and any proper subset I of vertices V ,

$$\sum_{i \in I} K_i(v) > - \sum_{(e,v) \in Lk(I)} (\pi - \Phi(e)) + 2\pi \chi(F_I), \tag{7}$$

where F_I is the set of all faces in M whose vertices are in I , $Lk(I)$ is the link of I being the set of pairs (e, v) of an edge e and a vertex v so that (1) the end points of e are not in I and (2) the vertex v is in I and (3) e and v form a triangle.

The following theorem is fundamental, showing that the map between the vertex radii Γ and the discrete curvature K is a homeomorphism; detailed proof can be found in [3].

Theorem 2. *If a discrete metric (M, Γ, Φ) induces a discrete curvature K , then K satisfies the Gauss–Bonnet equation (6) and the set of all linear inequalities (7). If M and Φ are given, K satisfies (6) and the set of all linear inequalities (7); then there exists a Γ unique up to scaling, such that K is induced by the metric (M, Γ, Φ) .*

The global surface parameterization problem can be reformulated as follows:

Global surface parameterization is meant to find a special metric, such that the curvatures are zero almost everywhere except at several singularities.

For example, conventional global conformal surface parameterization is done to compute a special metric on the mesh, such that at $|2g - 2|$ singularities, the curvatures are equal to -2π . The singularities are determined by the geometry of the surface, as shown in Fig. 2. The Ricci flow method allows the user to freely assign singularities for global parameterizations, as long as the target curvature satisfies the conditions in Theorem 2. Fig. 7 illustrates the relation between circle packing metric space and curvature space.

3.3. Discrete Ricci flow

One can assign a discrete Gaussian curvature \bar{K} for a weighted mesh (M, Φ) as long as \bar{K} satisfies the conditions in Theorem 2. Discrete Ricci flow is able to solve the vertex radii Γ . We use $e^u \Gamma$ to denote the conformal metric with vertex radius $e^{u_i} \gamma_i$ at vertex i . Similar to the continuous Ricci flow (1), we consider

Definition 3 (Discrete Ricci Flow). The discrete Ricci flow is defined as

$$\frac{du_i}{dt} = (\bar{K}_i - K_i), \tag{8}$$

where \bar{K}_i is the desired discrete Gaussian curvature at vertex v_i , under the constraint $\sum u_i = 0$ (equivalent to the area-preserving constraint).

Similar to continuous Ricci flow, it is proven that discrete Ricci flow also converges to this stable solution exponentially fast.

Definition 4 (Convergence). The solution to (8) is called *convergent* if

- (i) $\lim_{t \rightarrow \infty} K_i(t) = \bar{K}_i$ exists for all i ,
- (ii) $\lim_{t \rightarrow \infty} \gamma_i(t) = \bar{\gamma}_i \in \mathbb{R}^+$ exists for all i .

A convergent solution is called *convergent exponentially fast* if there are positive constants c_1, c_2 , so that for all time $t \geq 0$,

$$|K_i(t) - \bar{K}_i| \leq c_1 e^{-c_2 t},$$

and

$$|\gamma_i(t) - \bar{\gamma}_i| \leq c_1 e^{-c_2 t}.$$

The following theorem states that discrete Ricci flow is guaranteed to converge exponentially [3].

Theorem 5. *Suppose (M, Φ) is a closed weighted mesh. Given any initial circle-packing metric based on the weighted mesh, the solution to the discrete Ricci flow (8) in the Euclidean geometry with the given initial value always exists and converges exponentially fast. The solution converges to the metric $\Pi^{-1}(\bar{K})$.*

3.4. Conformality

In practice, it is highly desirable for the deformation to be conformal, namely, angle preserving. A conformal map transforms an infinitesimal circle to an infinitesimal circle, as shown in Fig. 1. Therefore conformal mapping only changes the radii γ in the circle packing metric (M, Φ, Γ) , and preserves the intersection angles Φ among the circles. It can be proven that continuous conformal mapping can be approximated with arbitrary accuracy by discrete maps using circle packing [20].

In graphics applications, the meshes are embedded in \mathbb{R}^3 , the metrics are induced from that of \mathbb{R}^3 . We can find the optimal weight Φ with initial circle radii Γ , such that the circle packing metric (M, Φ, Γ) is as close as possible to the Euclidean metric in the least square sense. Namely, we want to find (M, Φ, Γ) by minimizing the following functional

$$\min_{\Gamma, \Phi} \sum_{e_{ij} \in E} |l_{ij} - \bar{l}_{ij}|^2, \tag{9}$$

where \bar{l}_{ij} is the edge length of e_{ij} in \mathbb{R}^3 .

4. Affine atlas construction

In this section, we detail our algorithm for constructing the affine atlas by employing Ricci flow. The entire pipeline of the algorithm is illustrated in Fig. 9.

Step 1: Selecting singularities

We can select the singular vertices $\{v_1, v_2, \dots, v_k\}, k \geq 0$ anywhere on the mesh arbitrarily; then we assign the target curvature of the singular vertices such that

$$\sum_{i=1}^k \bar{K}(v_i) = 2\pi \chi,$$

where χ is the Euler number of the surface; the target curvature of other vertices are zero. Note that, there are several special cases that must be addressed.

- If the surface is a closed genus one mesh, then no singular vertex is needed.
- For a high genus mesh, we can select only one singular vertex and concentrate all curvatures on it.
- If the mesh is open, we can assign the target curvatures for all the interior vertices to be zero and assign the target curvatures for boundary vertices such that the total boundary curvatures equal $2\pi \chi$. By this way, all the non-zero curvature will be pushed to the boundary.

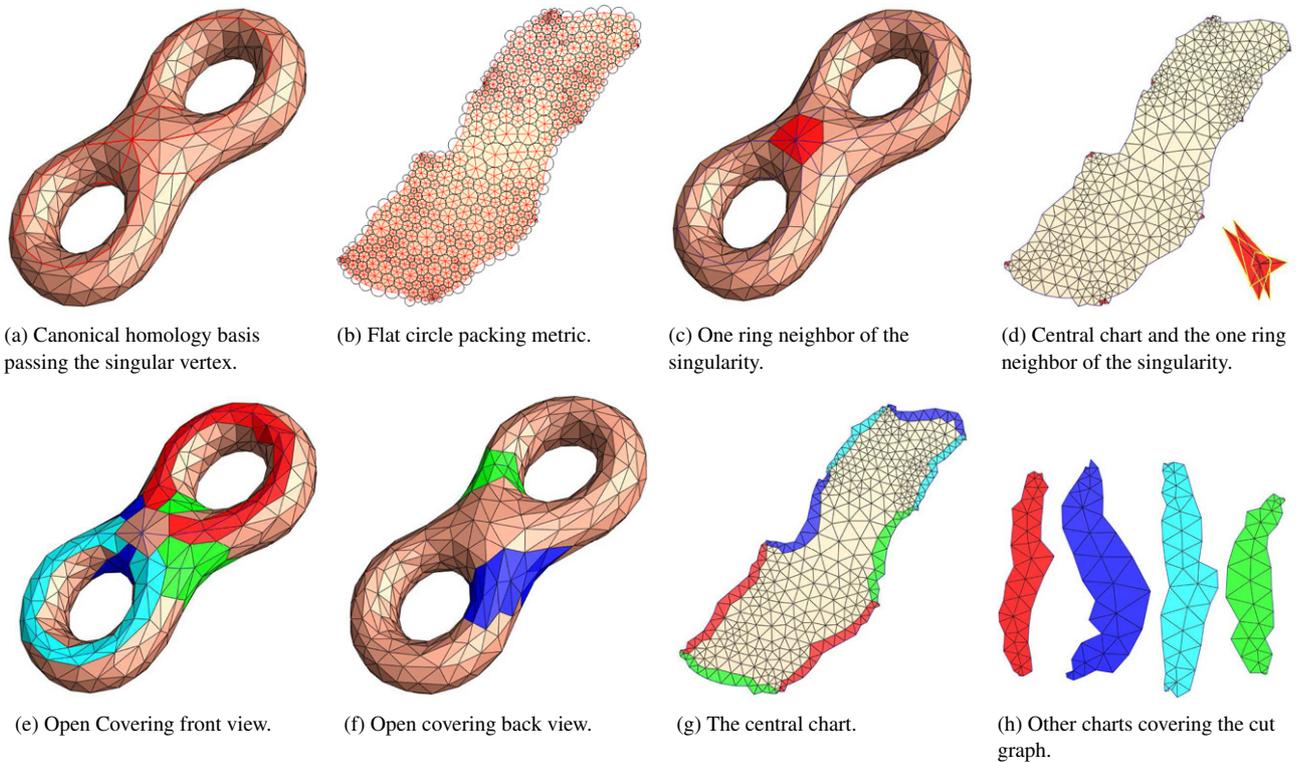


Fig. 8. Affine atlas automatically acquired by using Ricci Flow. First, the user selects one singular point as shown in (a). Then a cut graph is labelled either manually or automatically as the dark curves in (a), where the cut graph is a set of canonical homology basis curves passing through the singular vertex. Second, the flat circle packing metric is computed using Ricci flow, illustrated in (b). The flat metric induces a planar embedding. The entire surface is sliced open along the cut graph to form a topological disk (i.e. fundamental domain). The interior of the fundamental domain is bijectively mapped to the plane. The mapping of the one-ring neighbor of the singular vertex is not 1 to 1, but $2g - 1$ to 1 shown in (d). Other charts covering the cut graph are constructed as shown in (e) and (f), and their overlapping relation with the central chart is shown in (g) and (h) by encoded colors, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

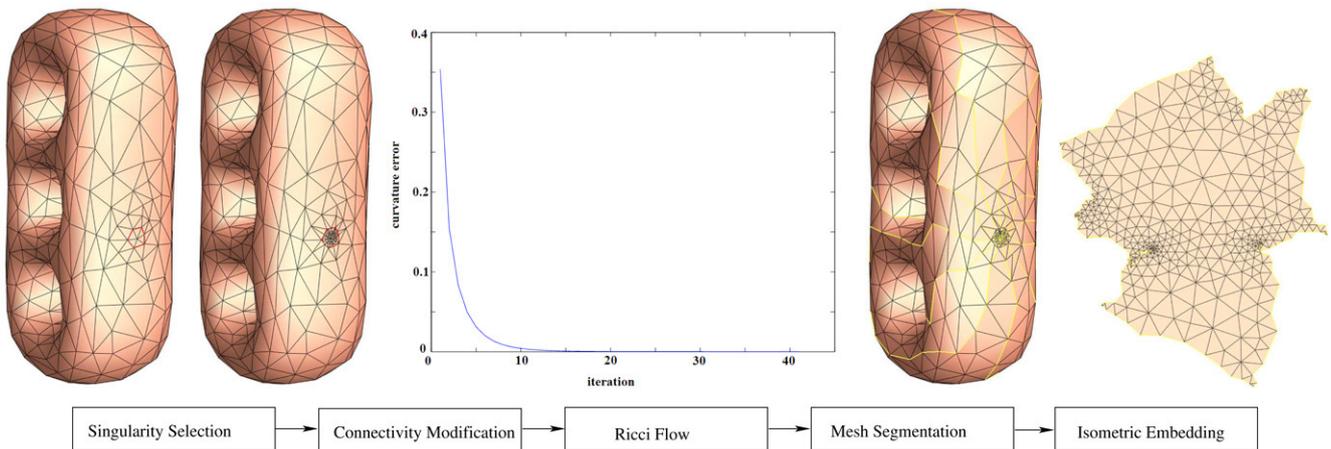


Fig. 9. The pipeline to compute the affine atlas using discrete Ricci flow.

Ricci flow only changes the vertex radii; therefore, the resulting metric is conformal to the original one, and no angle distortion will be introduced. But the area distortion is unavoidable. The uniformity of the parameterization varies drastically depending on the choice of singularities. Our selection is based on the minimal area distortion among all possible cases. For each case, we set one vertex as the single singularity.

The area distortion error is measured with the following equation:

$$\frac{\sum_{i=1}^n (-\log(S_i) - \log(s_i))^2}{n}$$

where S_i is the i -th triangle area in surface, s_i is its area in parametrization domain, and n is the total number of triangles.

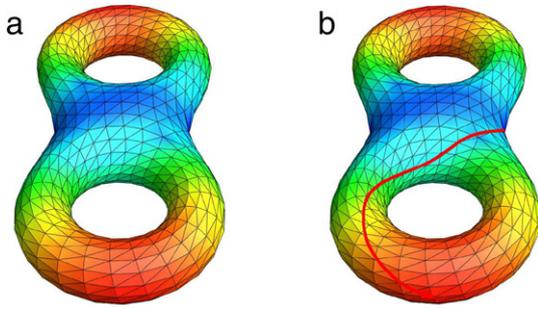


Fig. 10. (a) The distribution of the area distortion is color encoded. (b) The path of the arrow illustrate the rough behavior of the error distribution for the eight model. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Fig. 10(a) shows the distribution of the area distortion, blue area with low distortion, and red area with high distortion. In Fig. 10(b), the path of the arrow illustrates the rough behavior of the error distribution for the eight model.

When we sort the vertices with the area distortion errors, we may find the best position to put the singularity. Fig. 11 gives five different cases of setting singularity. From the left to the right, the area distortion is increasing, with the left most being the best parameterized and the right most being the worst parameterized.

Step 2: Modify local connectivity around the singular vertex

In order to determine the desired flat metric, the combinatorial constraints for the curvature (7) have to be satisfied. If both the initial curvature configuration and the target curvature configuration satisfy the constraints, any intermediate curvature configuration during Ricci flow will satisfy the constraints. Thus, it is enough to only consider the target curvature. If some singularities have high target curvature concentrations, we need to modify the local connectivity in their neighborhoods.

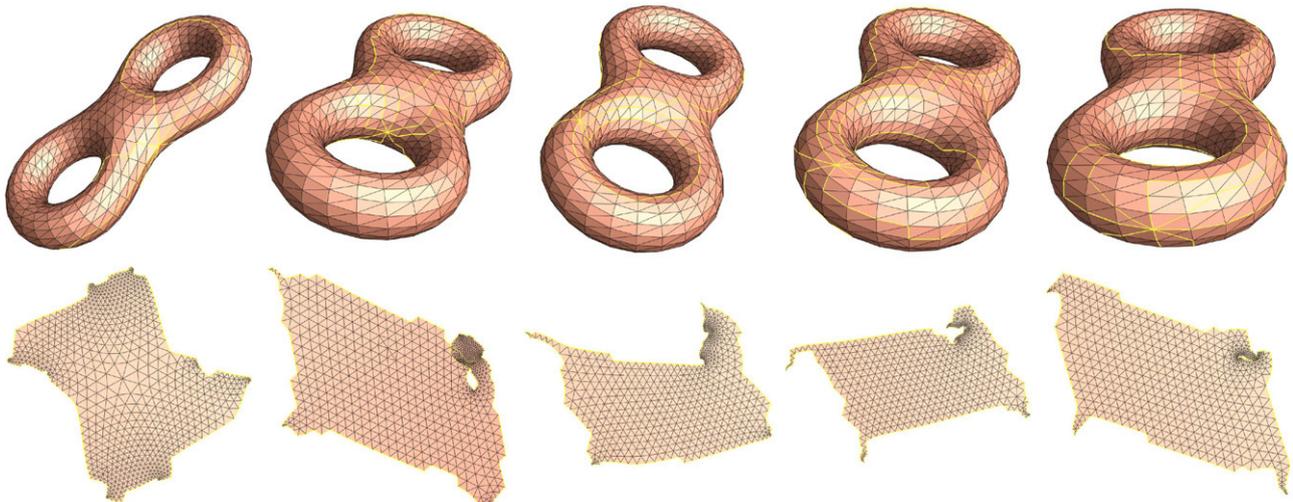


Fig. 11. The first row gives the different positions of singularities on the same model; The second row shows their corresponding flat metrics. From left to right, the area distortion increases.

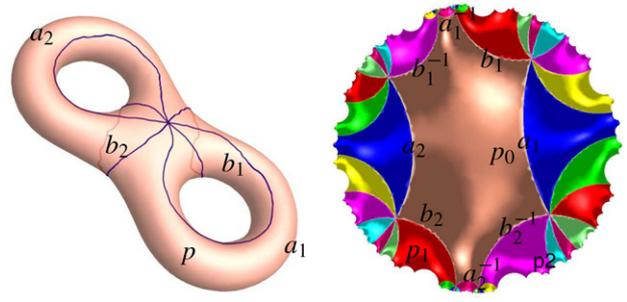


Fig. 12. A genus two surface with a set of canonical fundamental group generators $\{a_1, b_1, a_2, b_2\}$ is shown on the left. A finite portion of its universal covering space is shown on the right. Different fundamental domains are drawn in different colors. The boundary of each fundamental domain is the preimage of $a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1}$. The points $\{p_0, p_1, p_2\}$ are the preimages of p on the surface. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

In practice, we replace the combinatorial constraints Eq. (7) by a stronger one which is independent of the edge weight Φ ,

$$\sum_{i \in I} K_i(v) > - \sum_{(e,v) \in Lk(I)} \pi + 2\pi \chi(F_I).$$

We modify the connectivity around the extraordinary point and make the sampling in the neighborhood of the singularity much denser. We increase the connectivity of the singularity to no less than four times the genus, and vertices in the neighborhood have valence of about 6. This can be summarized as follows:

- (i) The valence of a singular vertex v is no less than $4 - \frac{2\bar{K}(v)}{\pi}$.
- (ii) For all the vertices in the first n ring neighbor of the singular vertex, their valences are no less than 6, n is a small integer. We set $n = 3$ in our experiments in this paper.

Fig. 15 demonstrates the step of the connectivity modification around the extraordinary point. This step can be easily done using the edge split operation in the half edge data structure.

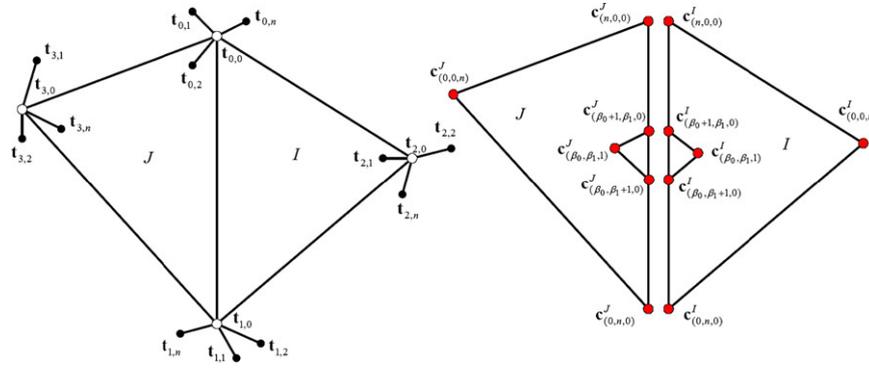


Fig. 13. The linear constraints in the least squares problem Eq. (14) for a cubic manifold triangular B -spline. (a) shows two adjacent triangles $I = \Delta(v_0, v_1, v_2)$ and $J = \Delta(v_3, v_1, v_0)$ and the knot configurations. (b) highlights six control points, three for each triangle. Note that the constraint for case $r = 0$ is equivalent to the shared control points, i.e., $c_{\beta_0+1,\beta_1,0}^J = c_{\beta_0+1,\beta_1,0}^I$ and $c_{\beta_0,\beta_1+1,0}^J = c_{\beta_0,\beta_1+1,0}^I$. The constraint for case $r = 1$ requires that the highlighted six control points to be coplanar.



Fig. 14. Examples of manifold splines with various extraordinary points. Rabbit: genus zero, two boundaries, no singularity; Vase: genus two, one singularity; Cup: genus two, one boundary, no singularity.

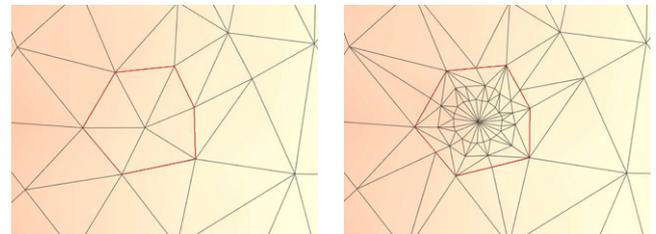
Step 3: Ricci flow

In order to compute the flat metric, we use Ricci flow to drive the mesh to deform in this order:

- (i) Set the initial value $u_i = 0$ for each vertex.
- (ii) Assign the weight for each edge and the radii for each vertex by minimizing the energy,

$$\min_{\Gamma, \Phi} \sum_{e_{ij} \in E} |l_{ij} - \bar{l}_{ij}|^2, \quad (10)$$

where \bar{l}_{ij} is the edge length of e_{ij} in \mathbb{R}^3 . Namely, the discrete metric (M, Φ, Γ) is consistent with the induced Euclidean metric on M .



(a) Original connectivity. (b) Modified connectivity.

Fig. 15. Modifying local connectivity around the extraordinary point. (a) The extraordinary point and its one-ring neighbors are marked in red. (b) Modifying local connectivity to satisfy the combinatorial constraint Eq. (7). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

- (iii) Update the vertex radius γ_i by $e^{u_i} \gamma_i$. Compute the current edge lengths l_{ij} 's using Eq. (3), corner angles, and discrete curvatures for each vertex using Eq. (4) or (5).
- (iv) Update vertex radii,

$$u_i + = \epsilon \times (\bar{K}_i - K_i), \quad (11)$$

where ϵ is a carefully selected step length. Note that if the step length is too small, the convergence is very slow. $\epsilon = 0.1$ is used in the experiments in this paper.

- (v) Normalize u_i , such that the summation of all u_i 's equals to zero.
- (vi) Check the deviation between K_i and \bar{K}_i , if the error is less than a predetermined threshold, the algorithm terminates. Otherwise, goto Step 2.

The algorithm will converge exponentially fast. In practice, the step length might be time-varying in order to improve the efficiency.

Step 4: Segmentation

Next, in order to construct the affine atlas, an open covering of the mesh needs to be built. The basic idea is to find a set of curves G such that the mesh M can be sliced open along the curves and form a topological disk. These kinds of curves form the *cut graph* as introduced in the work on geometry images [8].

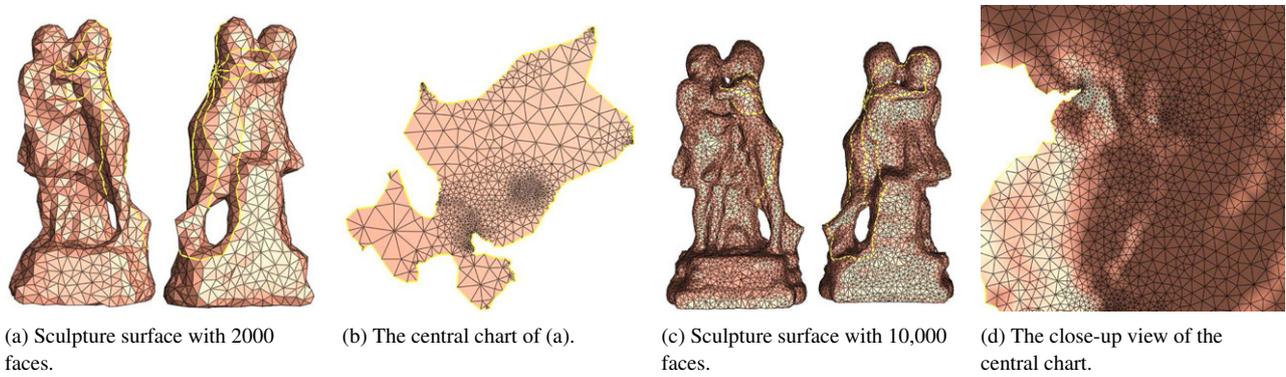


Fig. 16. Affine atlas using Ricci Flow for a genus 3 surface. The sculpture surface is of genus 3 with different resolutions, respectively, one is with 2000 faces as shown in (a), the other one is with 10,000 faces as shown in (c). A vertex is selected as the singularity, and a set of canonical homology basis curves passing through the singularity are drawn as the yellow curves in (a) and (c), respectively. By using Ricci flow, flat metrics on the meshes are computed such that all curvatures are concentrated on the singularities. The flattened fundamental domains are shown in (b) and (d), respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

If there is only one singularity p_0 , the cut graph can be constructed using a set of canonical homology bases passing through the singularity as introduced in [2]. The cut graph has one node and $2g$ edges. The edges can be labelled as $a_1, b_1, a_2, b_2, \dots, a_g, b_g$ (see Fig. 12).

Then the mesh is cut open along the cut graph to form a big chart \bar{M} . The boundary of \bar{M} has canonical form

$$\partial\bar{M} = a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1} \dots a_g b_g a_g^{-1} b_g^{-1}.$$

Each edge starts and ends at the singularity p_0 . We cover each edge s_k by a chart

$$U_k = \cup_{v_i \in s_k} N_i, \quad v_i \neq p_0, N_i = \cup f_{ijk},$$

where N_i represents the one ring neighbor of vertex v_i . The algorithm for computing an open covering of M is as follows:

- (i) Compute a cut graph G using a canonical homology basis.
- (ii) Slice the mesh along the cut graph to form a topological disk \bar{M} .
- (iii) For the edges of the cut graph, compute the union of one ring neighbors of all its interior vertices.
- (iv) The open covering of M is formed by \bar{M} and U_k ,

$$M/\{p_0\} \subset \bar{M} \cup_k U_k.$$

Step 5: Planar embedding

Because the curvature for each vertex is zero, the faces can be flattened one by one on the plane. The following algorithm describes the details on how to flatten an open set $U \subset M/\{p_0\}$. Let the desired parameterization is $\tau : U \rightarrow \mathbb{R}^2$,

- (i) Label all faces in U as non-processed ones. Meanwhile, label all vertices in U as non-processed.
- (ii) Select randomly a face $f_0 = [V_0, V_1, V_2]$ from U , label f_0 as processed, label all its vertices as processed. Assign $\tau(v_0) = (0, 0)$ and $\tau(v_1) = (l_{01}, 0)$. Compute $\tau(v_2)$ such that

$$|\tau(v_2) - \tau(v_0)| = l_{01}, |\tau(v_2) - \tau(v_1)| = l_{12}, \quad (12)$$

and

$$(\tau(v_1) - \tau(v_0)) \times (\tau(v_2) - \tau(v_0)) \cdot n > 0, \quad (13)$$

- (iii) Find all faces in U sharing an edge with f_0 , insert them to a face queue Q .
- (iv) If Q is empty, simply terminate. Otherwise, fetch the first face $f = [v_0, v_1, v_2]$ from Q , label f as processed.
- (v) If all vertices of f have been processed, go to Step 3. Otherwise, there must be only one vertex which has not been processed; assume it is v_2 , label v_2 as processed.
- (vi) Compute $\tau(v_2)$, such that both the distance condition (12) and orientation condition (13) are satisfied.
- (vii) Find all neighboring faces sharing an edge with f and they are not yet to be processed, add them to Q . Go to Step (iii).

In order to reduce the accumulation error, the parameterization can be further improved by minimizing the following functional:

$$\min_{\tau} \sum_{e_{ij}} (|\tau(v_i) - \tau(v_j)|^2 - l_{ij}^2)^2.$$

The purpose of the above functional is to find a valid embedding such that the distortion between the edge length in the parametric domain and the Ricci flow output is minimal. In practice, this step is usually unnecessary if singularities are carefully chosen to spread out the surface and the curvature for each of them is not extremely high.

Figs. 3, 8 and 16 demonstrate the affine atlas for surfaces from genus one, two and three, respectively.

5. Manifold spline construction

After the affine atlases are constructed in the previous section, this section first briefly summarizes the theory of manifold splines, and then presents our experimental results.

Suppose M is a mesh with the one ring neighbors of the singular vertices removed. $\{(U_i, \tau_i)\}$ is an affine atlas, where U_i is a topological disk comprised of a set of faces of M , $\tau_i : U_i \rightarrow \mathbb{R}^2$ maps U_i onto the plane, namely, (U_i, τ_i) forms a local coordinate chart. The chart transition functions $\tau_{ij} : \tau_i(U_i \cap U_j) \rightarrow \tau_j(U_i \cap U_j)$ is a rigid-body motion in \mathbb{R}^2 .

A manifold spline is defined on the mesh $\mathbf{F} : M \rightarrow \mathbb{R}^3$, such that

- The local representations of manifold splines, $\mathbf{F} \circ \tau_i^{-1} : \tau_i(U_i) \rightarrow \mathbb{R}^3$, are commonly used spline schemes with planar parameter domain.
- The evaluation of manifold splines is independent of the choice of local parameter charts,

$$\mathbf{F} \circ \tau_i^{-1} = \mathbf{F} \circ \tau_j^{-1} \circ \tau_{ij}$$

In our current implementation for this paper, we use triangular B -splines [4], because it has no restrictions on the connectivity of the mesh and it can represent any polynomials defined over planar [10], spherical [11] and manifold domain [8]. We have implemented our own version based on a generic half edge mesh library as in [15], while adding the edge lengths, vertex radii, and curvatures as the new attributes for the underlying mesh.

In our prototype software system, we have tested several meshes of genus from zero to three. In this paper, we choose the manifold triangular B -spline because of its flexibility in domain construction. This method can be also applied to other manifold splines, such as T-splines [14] and Powell-Sabin splines [12].

Given a domain manifold M , a manifold triangular B -spline surface is defined as follows:

$$\mathbf{F}(\mathbf{u}) = \sum_I \sum_{|\beta|=n} \mathbf{c}_{I,\beta} N_{I,\beta}(\tau_I(\mathbf{u})), \mathbf{u} \in M,$$

where I is the triangle index and $\beta = (\beta_0, \beta_1, \beta_2)$ is the 3-tuple used to label the control points and knots. The algorithm for constructing the manifold triangular B -spline is as follows:

- The initial control points $\mathbf{c}_{I,\beta}$ are chosen by uniformly subdivided the domain manifold M according to the user-specified degree n . Each domain triangle is associated with $(n+1)(n+2)/2$ control points.
- To compute the optimal control points $\tilde{\mathbf{c}}_{I,\beta}$, we solve the following linear constrained least square problem:

$$\min_{\tilde{\mathbf{c}}} \sum_I \sum_{|\beta|=n} \|\tilde{\mathbf{c}}_{I,\beta} - \mathbf{c}_{I,\beta}\|^2 \quad (14)$$

subject to $\tilde{\mathbf{c}}_{I,\beta} = f^J(V_\beta^I), \forall I, \forall \beta, |\beta| = n, \beta_2 \leq r$

where $I = \Delta(v_0, v_1, v_2)$ and $J = \Delta(v_3, v_1, v_0)$ are adjacent triangles with common edge $\{v_0, v_1\}$, $V_\beta^I = \{\mathbf{t}_{0,0}^I, \dots, \mathbf{t}_{0,\beta_0-1}^I, \dots, \mathbf{t}_{2,0}^I, \dots, \mathbf{t}_{2,\beta_2-1}^I\}$ and $\mathbf{t}_{i,j}^I$ are the knots for triangle I .

Note that the initial manifold triangular B -spline surfaces acquired by step 1 usually have very bad curvature distribution, especially along the edges of the domain triangles. The purpose of step 2 is to fair the spline surface by modifying the control points. In the objective function Eq. (14), we minimize the squared distance between the control points of the original and the new spline surfaces, which implies a minimal change of the shape. In the constraints, we use an integer r , $0 \leq r \leq n-1$, to control the fairness of the spline surface. The bigger the value r , the more faired the surface we obtain. In our experiments, we can get visually pleasing surfaces with $r = 1$ for cubic splines or $r = 2$ for splines of degree 5 or above. Fig. 13 illustrates the case $r = 1$. For the detailed information about spline fairing, please refer to [13].

Fig. 8 shows the flat circle packing metric of a genus two surface and its affine atlas. Fig. 9 demonstrates the process of using Ricci flow on how to compute the affine atlas. The sculpture surface in Fig. 16 is of genus three with different resolutions. The singular vertex and the cut graph are explicitly shown in this figure. The affine atlases are also highlighted in the figure.

All the examples of manifold triangular B -splines are shown in Figs. 14, 17 and 18. Table 1 shows the statistics of the test cases. As shown in this table, our algorithms for constructing the affine atlas and manifold triangular B -splines are extremely fast, i.e., within only a few seconds.

Note that genus 1 surfaces, such as Rockerarm (closed) and Hypersheet (open), do not have singularities. There is no singularity for the cup model (genus 2 open surface) as all the non-zero curvatures are pushed to the boundary. For the Rabbit model, we introduce two cuts, one on the top, the other on the bottom, and then apply the double covering to convert it into a closed genus one surface. Thus, no singularity exists for the Rabbit model. For the Bunny model, we specify the singularity on the head. For other genus 0 surfaces with large extrusion parts, such as Horse and Camel, we choose multiple singularities on their legs to reduce the area distortion of the affine atlas. For the high genus closed models, such as Eight, Vase and 3-hole torus, only a single extraordinary point is specified. As we remove the extraordinary points and their one-ring neighbors in spline construction, there are holes in the resulting splines. In the postprocessing step, we use planar triangular B -splines to fill these holes with G^1 continuities along the boundaries.

6. Discussion

This section discusses several implementation related issues.

Conformal structure vs affine structure. It is proven that conformal structure induces affine structure but not vice versa [8]. Therefore, conformal structure is in some sense stronger than affine structure. In [8], Gu et al. showed a method to construct manifold triangular B -splines using conformal structure. In fact, manifold splines are solely defined using affine structures, since all the popular planar splines such as NURBS, Bèzier splines, triangular B -splines are parametric affine invariant. Thus, from the spline construction point of view, one can totally ignore the conformal constraints. The affine atlas shown in Fig. 9 is computed by ignoring the conformal constraints, i.e., we simply skip Step 3 (ii) and assign the initial radius of each vertex to be 1. Therefore, the angle Φ_{ij} in the circle packing metric is zero for each edge e_{ij} . Since discrete Ricci flow preserves the Φ_{ij} , the triangles in the isometric embedding tend to be equilateral triangles. We should also point out that remeshing of the input domain mesh is not necessary, except that the valance of the singular point needs to be increased.

Number of singularities. The intrinsic connection between manifold splines and polar forms results from the affine structure and affine atlas. To make these geometric structures computational tractable in shape modeling applications, we resort to the powerful tool of global parameterization over arbitrary manifold domain. It is known that global



Fig. 17. Multiple singularities. Within our framework, the user can also specify multiple singularities on the models. The horse model has five boundaries (four on the feet and one on the mouth) and two singularities. The camel model has five boundaries and one singularity. The target geodesic curvature of the boundary vertices is zero. The holes and singularities are filled using minimal surfaces.

parameterization often suffers from severe area distortion. The quality of the global parameterization is determined by many factors, such as the connectivity of the mesh, the weights on edges, the positions and curvatures of the singularities. It is technically challenging to optimize these factors towards the quality improvement of global parameterization.

Section 4 presents a brutal force method to locate the singularity such that the area distortion of the affine atlas is minimal. Usually, the larger the number of singularities, the less area distortion in the affine atlas. In practice, the singular points should be chosen on parts with large extrusions, e.g., the feet of the horse model in Fig. 17.

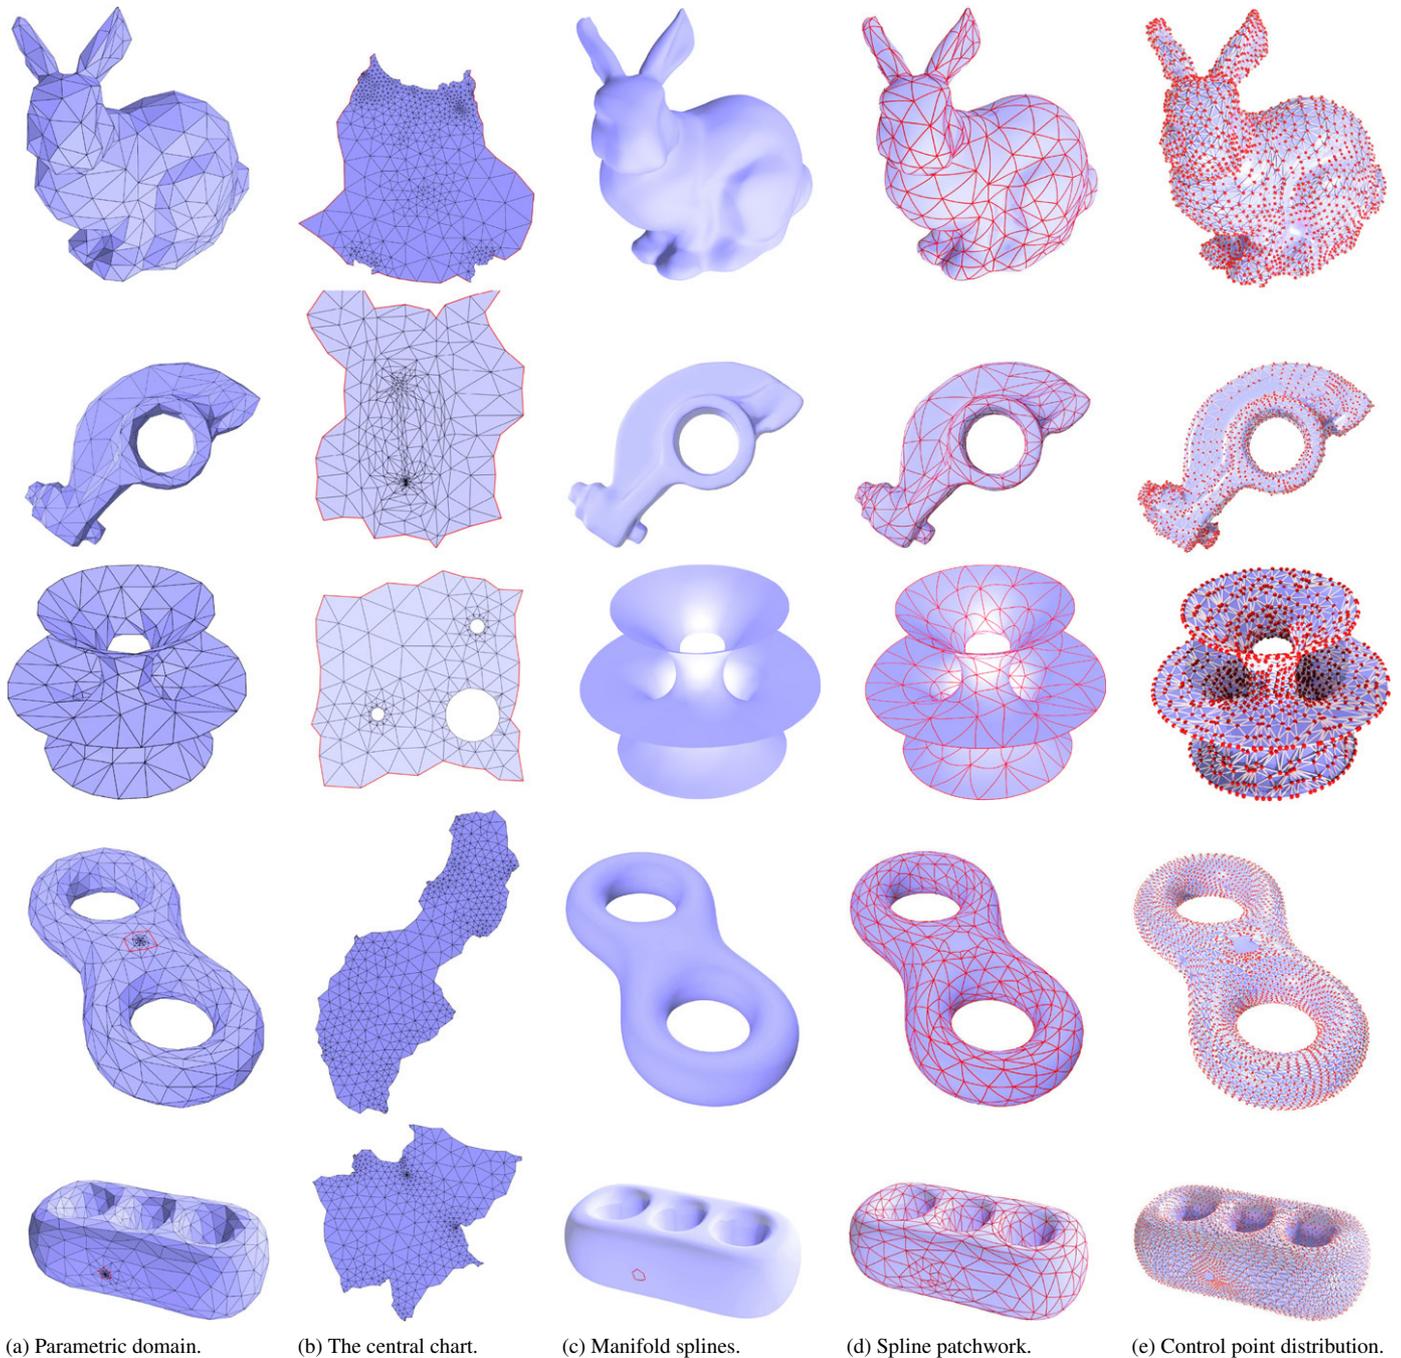


Fig. 18. Examples of manifold triangular B -splines. The affine atlases are computed using Ricci flow under free boundary conditions. The transition function is a combination of translation and rotation. The red curves on the spline surfaces (see (d)) highlight the triangular patchwork. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Besides this work which minimizes the *number of singularities*, we developed polycube T-splines which aim to minimize the *area distortion* in the affine atlas [26]. However, the price to pay for the lower area distortion is the significantly larger number of extraordinary points. Therefore, in real applications, it is the user's call to make the balance.

7. Conclusion and future work

This paper has developed an efficient and rigorous algorithm for constructing a manifold spline surface of complicated

topology and complex geometry with single extraordinary point, which has already reached the theoretical lower bound of the number of singularities. The uniqueness of this construction algorithm for manifold splines is that, it is solely based on a simple and powerful computational tool: Ricci flow. From the mathematical point of view, Ricci flow has substantial relevance to the curvature flow method in differential geometry. For example, Ricci flow can conformally deform the metric to induce any prescribed curvature.

Table 1
Statistics of test examples

Object	g	N_f	N_b	N_s	N_c	T_{ricci} (s)	T_{spline} (s)	n	Continuity
Bunny	0	621	3	1	2827	6	3	3	C^2
Rabbit	0	1038	2	0	4698	12	8	3	C^2
Horse	0	4002	5	2	18074	26	18	3	C^2
Camel	0	2958	5	1	13380	18	12	3	C^2
Rockerarm	1	614	0	0	7675	2	5	3	C^2
Hypersheet	1	300	3	0	1446	5	2	3	C^2
Eight	2	806	0	1	3644	4	3	3	C^2
Vase	2	1480	0	1	6666	15	10	3	C^2
Cup	2	1929	1	0	15515	20	18	4	C^3
3-hole torus	3	878	0	1	3955	8	4	3	C^2

g , genus of domain manifold M ; N_f , # of faces in M ; N_b , # of boundaries in M ; N_s , # of singularities; N_c , # of control points; T_{ricci} , time for computing the discrete Ricci flow and isometric embedding (Step 3 and 5 in Section 4); T_{spline} , time for the spline construction; n , degree of splines. Note that time measures are in seconds.

Current manifold splines are essentially founded upon the natural integration of the affine atlases for domain manifolds and polar forms used to define conventional spline surfaces over any planar domain. Due to their topological obstruction, general high genus surfaces admit neither a flat metric nor an affine atlas. Therefore, ideally the most natural spline solutions for high genus manifolds should not depend on their affine structures. In the future, we shall investigate different spline schemes which are not based on the affine structure of the underlying surface domain.

Acknowledgements

The authors would like to thank Junho Kim for his help on preparing Figs. 10 and 11. Xianfeng Gu and Miao Jin are partially supported by NSF CCF-0448399, DMS-0528363, DMS-0626223, IIS-0713145, and NSF China 60628202. Ying He is partially supported by NTU/SUG 19/06 and AcRF RG69/07. Hong Qin is partially supported by NSF IIS-0326388, IIS-0601168, and IIS-0710819.

References

- [1] Bobenko AI, Springborn BA. Variational principles for circle patterns and Koebe's theorem. *Trans Amer Math Soc* 2004;356:659–89.
- [2] Carner Christopher, Jin Miao, Gu Xianfeng, Qin Hong. Topology-driven surface mappings with robust feature alignment. In: *Proceedings of IEEE visualization*. 2005. p. 543–50.
- [3] Chow Bennett, Luo Feng. Combinatorial Ricci flows on surfaces. *J Differential Geom* 2003;63(1):97–129.
- [4] Dahmen W, Micchelli CA, Seidel H-P. Blossoming begets B -spline bases built better by B -patches. *Math Comput* 1992;59(199):97–115.
- [5] Floater MS, Hormann K. Surface Parameterization: A Tutorial and Survey. In: *Advances in Multiresolution for Geometric Modelling*. 2005. p. 157–86.
- [6] Grimm Cindy, Hughes John F. Modeling surfaces of arbitrary topology using manifolds. In: *Proceedings of SIGGRAPH*. 1995. p. 359–68.
- [7] Gu Xianfeng, Yau Shing-Tung. Global conformal surface parameterization. In: *SGP*. 2003. p. 127–37.
- [8] Gu Xianfeng, He Ying, Qin Hong. Manifold splines. *Graph Models* 2006; 68(3):237–54.
- [9] Hamilton RS. The Ricci flow on surfaces. In: *Mathematics and general relativity* (Santa Cruz, CA, 1986). *Contemp Math* 1998;71: 237–262.
- [10] He Ying, Qin Hong. Surface reconstruction with triangular B -splines. In: *Proceedings of geometric modeling and processing*. 2004. p. 47–58.
- [11] He Ying, Gu Xianfeng, Qin Hong. Rational spherical spline for genus zero shape modeling. In: *Proceedings of IEEE international conference on shape modeling and applications*. 2005. p. 82–91.
- [12] He Ying, Jin Miao, Gu Xianfeng, Qin Hong. A C^1 globally interpolatory spline of arbitrary topology. *LNCS*, vol. 3752. 2005. p. 295–306.
- [13] He Ying, Gu Xianfeng, Qin Hong. Automatic shape control of triangular B -splines of arbitrary topology. *J Comput Sci Tech* 2006;21(2): 232–237.
- [14] He Ying, Wang Kexiang, Wang Hongyu, Gu Xianfeng, Qin Hong. Manifold T-spline. *LNCS*, vol. 4077. 2006. p. 409–22.
- [15] Hoppe Hugues. Efficient implementation of progressive meshes. *Comput Graph* 1998;22(1):27–36.
- [16] Jin Miao, Luo Feng, Gu Xianfeng. Computing surface hyperbolic structure and real projective structure. In: *ACM symposium on solid and physics modeling*. 2006.
- [17] Kharevych Lilya, Springborn Boris, Schröder Peter. Cone singularities to the rescue: Mitigating area distortion in discrete conformal map. In: *ACM SIGGRAPH/Eurographics symposium on geometry processing*. 2005.
- [18] Ni Xinlai, Garland Michael, Hart JC. Fair morse functions for extracting the topological structure of a surface mesh. *ACM Trans Graph* 2004; 23(3):613–22.
- [19] Ray Nicolas, Li WC, Levy Bruno, Sheffer Alla, Alliez Pierre. Periodic global parameterization. *ACM Trans Graph* 2006.
- [20] Rodin Burt, Sullivan Dennis. The convergence of circle packings to the Riemann mapping. *J Differential Geom* 1987;26(2):349–60.
- [21] Seidel Hans-Peter. An introduction to polar forms. *IEEE Comput Graph Appl* 1993;13(1):38–46.
- [22] Sheffer Alla, Praun Emil, Rose Kenneth. Mesh parameterization methods and their applications. *Foundations and Trends® in Computer Graphics and Vision* 2006;2(2).
- [23] Stephenson Kenneth. *Introduction to circle packing: The theory of discrete analytic functions*. Cambridge University Press; 2005.
- [24] Thurston W. Hyperbolic geometry and 3-manifolds, Low-dimensional topology. *Lecture Note Ser.*, vol. 48. London Math. Soc.; 1982. p. 9–25.
- [25] Ying Lexing, Zorin Denis. A simple manifold-based construction of surfaces of arbitrary smoothness. *ACM Trans Graph* 2004;23(3): 271–275.
- [26] Wang Hongyu, He Ying, Li Xin, Gu Xianfeng, Qin Hong. Polycube splines. In: *Proceedings of ACM symposium on solid and physical modeling*. 2007. p. 241–51.