

Conformal geometry of simplicial surfaces

Keenan Crane

1. Overview

What information about a surface is encoded by *angles*, but not *lengths*? This question encapsulates the basic viewpoint of *conformal geometry*, which studies *holomorphic* or (loosely speaking) *angle- and orientation-preserving* maps between manifolds. In the discrete setting, however, the idea of literally preserving angles leads to an interpretation of conformal geometry that is far too *rigid*: the space of discrete solutions looks far more restricted than the space of smooth solutions it hopes to model (Figure 1). Instead, one must consider how several equivalent points of view in the smooth setting lead to a number of *inequivalent* treatments of conformal geometry in the discrete setting. This activity has recently culminated in a complete *discrete uniformization theorem* for polyhedral surfaces, which beautifully mirrors the classic uniformization theorem for Riemann surfaces.

Why study discrete conformal geometry? From an analytic point of view, *smooth* conformal maps provide a strong notion of regularity, since they are complex analytic and hence have derivatives of all orders. In the discrete setting one therefore obtains one possible notion of what it means for a discrete (*e.g.*, simplicial) map to be “regular,” even though no derivatives are available in the classical sense. From a topological point of view, conformal maps help one define canonical mappings between spaces—for instance, uniformization provides an explicit map between any two conformally equivalent surfaces, by passing through a canonical domain of constant curvature. Likewise, discrete conformal geometry can be used as a starting point for constructing canonical maps between simplicial surfaces, even when they do not have compatible triangulations [HK15, BCK18]. In applications, discrete conformal geometry has become a powerful tool for digital geometry processing algorithms, since many tasks ultimately boil down to solving sparse linear systems or easy convex optimization problems. Discrete conformal maps have hence become essential for tasks ranging from regular surface remeshing [CZ17], to machine learning [MGA⁺17], to digital fabrication [KCD⁺16, SC18]. Discrete conformal geometry also arises in the context of statistical mechanics, where theories based on

1991 *Mathematics Subject Classification*. Primary 52C26; Secondary 52-02.

Key words and phrases. Discrete differential geometry, conformal geometry.

The author was supported in part by NSF Award #1717268 and a Packard Fellowship.




FIGURE 1. In the smooth setting, conformal maps can be characterized as maps that preserve angles and orientation. In the discrete setting, naïve applications of this characterization lead to definitions that are far too rigid: for instance, plane triangulations that share interior angles are identical up to a global similarity transformation.

discrete harmonic functions [Smi10] and circle patterns [Lis17, KLRR18] appear to be quite natural.

Discretization of conformal maps is an excellent example of “The Game” played in discrete differential geometry [CW17], since there are a large number of different (yet equivalent) characterizations of conformal maps in the smooth setting, each of which leads to a distinct starting point for a discrete definition. Consider for instance a smooth, nondegenerate, and orientation-preserving map f from a disk-like surface M with Riemannian metric g to the flat complex plane \mathbb{C} . A reasonably comprehensive list of ways to assert that f is *conformal* includes:

- (1) (**ANGLES**) Preservation of angles—at every point $p \in M$ the angle between any two tangent vectors $X, Y \in T_p M$ is the same as the angle between their images $df_p(X), df_p(Y)$ in the plane.
- (2) (**CIRCLES**) Preservation of circles—the image of any geodesic circle of radius ε approaches a Euclidean circle as ε goes to zero.
- (3) (**ANALYTIC**) The map f satisfies the *Cauchy-Riemann equation*— $df(\mathcal{J}X) = \iota df(X)$ for all vector fields X , where ι is the complex unit and $\mathcal{J}_p : T_p M \rightarrow T_p M$ is the *linear complex structure* expressing a rotation by $\pi/2$ in each tangent space ($\mathcal{J}_p^2 = -\text{id}$).
- (4) (**METRIC**) Conformal equivalence of metrics—the metric g on M is related to the induced metric $\tilde{g} := df \otimes df$ via a positive scaling $\tilde{g} = e^{2u}g$, where $u : M \rightarrow \mathbb{R}_{>0}$ is called the *log conformal factor*.
- (5) (**CONJUGATE**) Conjugate harmonic functions—the map f can be expressed as $f = a + bi$, where $a, b : M \rightarrow \mathbb{R}$ are harmonic functions with orthogonal gradients of equal magnitude ($\nabla b = \mathcal{J}\nabla a$).
- (6) (**DIRICHLET**) Critical points of Dirichlet energy—among all homeomorphisms from M to $f(M)$, the map f is a critical point of the *Dirichlet energy* $E_D(f) := \int_M |df|^2 dA$, where dA is the area measure on (M, g) .
- (7) (**HODGE**) Hodge duality—the Hodge star on differential 1-forms induced by f is the same as the pushforward under f of the 1-form Hodge star on the original domain.

To someone familiar with smooth conformal geometry this list may seem redundant, since in most cases the conceptual leap between one characterization and another is very small. Yet these minor shifts in perspective often lead to substantially different interpretations in the discrete setting. In almost all cases, the “discrete setting” means we are considering a domain M encoded as either

- (1) a triangulated surface, or
- (2) a quadrilateral net.

These notes focus primarily on triangulations; see [BG17] and references therein for an overview of the quadrilateral case. Our aim is to give a broad overview of some of the highlights of discrete conformal geometry; many details have therefore been omitted (though see the bibliography for a detailed list of references). The notes are organized into four parts:

- **Part I** explores how a naïve treatment of discretization (*e.g.*, preservation of angles in the triangulation, or direct discretization of the Cauchy-Riemann equation) leads to definitions of discrete conformal maps that are excessively rigid, or agree with the smooth theory only in the limit of refinement.
- **Part II** shows the first glimpse of real “discrete” theories based on preservation of circles (*circle packings* and *circle patterns*). This perspective captures many important features of conformal geometry, and very nearly provides a complete theory of discrete uniformization for general triangulations.
- **Part III** considers the perspective of conformally equivalent discrete metrics; this is to date the most satisfactory theory of conformal maps for triangulated surfaces, since it comes with a complete uniformization theorem.
- **Part IV** makes the connection between discrete conformal mapping problems and realization problems for ideal hyperbolic polyhedra, which also illuminates the relationship between theories based on circles (Part II) and conformally equivalent metrics (Part III).

Importantly, the focus in these notes on *structure preservation* should not be confused with a value judgement on *utility*: numerical schemes that do not exactly capture smooth structure may nonetheless be perfectly suitable for practical computation (especially on fine tessellations), and are often less computationally demanding than those that furnish an exact discrete theory. The same trade off can be found throughout discrete differential geometry: exact structure preservation often comes at significant computational cost relative to cheaper numerical alternatives. A negative interpretation is that one is therefore stuck between fast but inexact numerical schemes, and those that are “exact” but slow. A more mature point of view is that the two sets of tools are complementary: fast numerical methods help to initialize or approximate intermediate computations needed for exact, structure-preserving schemes, which in turn provide valuable guarantees about the behavior of algorithms. Beyond computation, discrete conformal geometry helps to bridge several areas of mathematics, including some rather remarkable connections between geometry, analysis, and combinatorics, as well as Euclidean and hyperbolic geometry. It also highlights a central thesis of discrete differential geometry, namely that the most important features of geometry are not inherently smooth nor discrete, but (as we will see) can be faithfully described in either language.




FIGURE 2. Topologically, we consider surfaces obtained by gluing triangles along edges. A torus with one vertex (top), cone with two vertices (middle), and sphere with three vertices (bottom) can each be expressed as a Δ -complex, but not as a simplicial complex.

2. Preliminaries

Throughout we consider a compact connected surface M with Riemannian metric g . In Part I we will mainly consider a differentiable map f from a disk-like domain M to the complex plane \mathbb{C} (with its usual Euclidean metric); this setup allows us to focus on the local picture, and avoid issues of global topology or complex structure. As discussed in Section 1, there are many equivalent ways to express that f is *conformal*; we therefore refrain from choosing a canonical definition, and instead discuss each characterization in turn. The differential $df_p : T_p \rightarrow T_{f(p)}\mathbb{C}$ of the map f expresses how tangent vectors on M are pushed forward to \mathbb{C} . We will generally assume that f is an *immersion*, meaning that its differential df is nondegenerate, *i.e.*, at each point $p \in M$, $df_p(X) = 0$ if and only if $X = 0$. In particular, every conformal map is an immersion, whereas holomorphic maps can have isolated points where the differential fails to be injective. Finally, the *linear complex structure* associated with a surface (M, g) is a tangent space automorphism $\mathcal{J} : TM \rightarrow TM$ such that in each tangent space $T_p M$, (i) $\mathcal{J}_p^2 = -\text{id}$ (where id denotes the identity map), and (ii) $g_p(X, \mathcal{J}_p X) = 0$; intuitively, \mathcal{J} defines a quarter-rotation in each tangent space compatible with the metric g (in analogy with the imaginary unit i in the complex plane), and is hence determined by the metric up to a global choice of orientation.

In the discrete setting, we consider surfaces that can be obtained by gluing together Euclidean triangles—the prototypical example is a Euclidean polyhedron (such as the icosahedron), though we need not restrict ourselves to surfaces embedded

in \mathbb{R}^n . Usually it will be sufficient to consider simplicial complexes, though for the general theory of discrete conformal equivalence (Part III) we must consider more general triangulations—here it will be sufficient to work with Δ -complexes, as defined by Hatcher [Hat02, Section 2.1]. Figure 2 gives several examples of triangulations that are not simplicial but are expressible as Δ -complexes. Throughout, we will use the term *discrete surface* to mean a topological 2-manifold triangulated by either a simplicial or Δ -complex $M = (V, E, F)$, where V , E , and F denote the vertices (0-cells), edges (1-cells), and faces (2-cells), *resp.* We will use a $(k + 1)$ -tuple of vertex indices to specify a k -simplex; for instance, ij is an edge with vertices $i, j \in V$ and ijk is a face with vertices $i, j, k \in V$. (This convention is an abuse of notation in the case of a Δ -complex, but the meaning will nonetheless be clear.) If a discrete surface has finitely many triangles, we say that it is *finite* (and hence compact). The *star* $\text{St}(i)$ of a vertex $i \in V$ is the subcomplex of M comprised of all simplices that contain i .

The intrinsic geometry of a discrete surface $M = (V, E, F)$ can be expressed via a *discrete metric*, *i.e.*, an assignment of positive edge lengths $\ell : E \rightarrow \mathbb{R}_{>0}$ that strictly satisfy the triangle inequality in each face:

$$\ell_{ij} + \ell_{jk} > \ell_{ki} \quad \forall ijk \in F.$$

These edge lengths naturally define a piecewise Euclidean metric on M , obtained by constructing disjoint Euclidean triangles with the prescribed lengths and gluing them along shared edges. The resulting space has a singular Riemannian metric g that is Euclidean away from vertices, and “cone-like” in a small neighborhood around each vertex $i \in V$. Any such metric g is called a (*polyhedral*) *cone metric*; a more formal definition is given by Troyanov [Tro91, p. 4].

Any discrete metric determines interior angles θ_i^{jk} at each vertex i of each triangle ijk (*e.g.*, via the law of sines). Letting

$$\Theta_i := \sum_{ijk \in F} \theta_i^{jk}$$

be the total angle around vertex i , the *cone angle*

$$(2.1) \quad \Omega_i := 2\pi - \Theta_i$$

measures how close the vertex is to being Euclidean, providing a discrete analogue for Gaussian curvature. Intuitively, the intrinsic geometry looks like a circular Euclidean cone, or more generally, a circular wedge of the Euclidean plane glued together at opposite edges (Figure 3). Likewise, for any vertex i on the domain boundary, the angle $\kappa_i := \pi - \sum_{ijk \in F} \theta_i^{jk}$ provides a discrete analogue for geodesic curvature. Together, Ω and κ satisfy a discrete Gauss-Bonnet theorem: $\sum_i \Omega_{i \in M} + \sum_{i \in \partial M} \kappa_i = 2\pi\chi(M)$, where $\chi(M) := |V| - |E| + |F|$ is the Euler characteristic of M .

For an embedded surface, the extrinsic geometry can be expressed as a piecewise linear (or more properly, simplexwise affine) map interpolating given vertex coordinates $f : V \rightarrow \mathbb{R}^n$. Any such map is a *discrete immersion* if it is locally injective, or equivalently, if for each vertex $i \in V$ the restriction of f to $\text{St}(i)$ is embedded [Cer96, Lemma 2.2]. The condition that f be a (discrete) immersion not only excludes vanishing angles, zero-length edges, and zero-area triangles, but also avoids discrete branch points (see Figure 4). It would therefore be reasonable to require any definition of a discrete conformal map to include the condition that f is a discrete

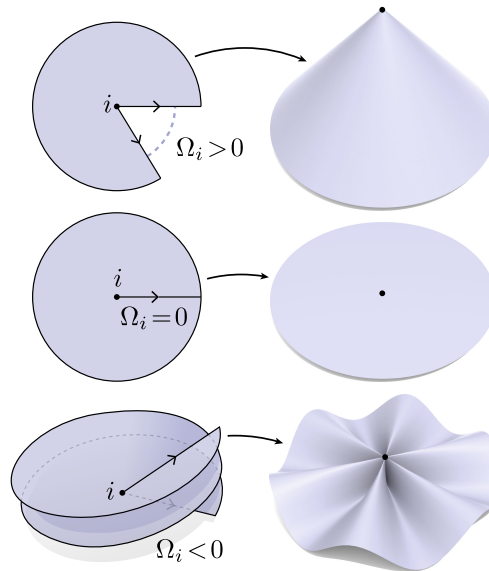


FIGURE 3. A cone metric is flat away from isolated *cone points* $i \in V$, where it looks like a wedge of paper (possibly multiply-covered) identified along the two marked edges. The angle deficit or excess Ω_i determines the curvature.




FIGURE 4. A discrete branch point, where a piecewise linear map fails to be a discrete immersion.

immersion; a discrete holomorphic map should likewise be noninjective only at vertices. Any discrete immersion f induces a corresponding discrete metric

$$\ell_{ij} := |f_j - f_i|$$

(where $|\cdot|$ denotes the Euclidean norm); note that unlike the smooth setting this metric is nondegenerate even in the presence of branch points.

Since conformal maps depend only on the intrinsic geometry of a surface, the *edges* of a polyhedron have no geometric significance—even when the metric arises from an embedding in \mathbb{R}^n . For instance, consider an embedded polyhedron with two triangles ijk , jil sharing an edge ij (Figure 5). Since this triangle pair is isometric to a Euclidean quadrilateral (with two possible diagonals), an intrinsic observer walking along the surface has no way to know when they are crossing the extrinsic edge. Likewise, two discrete surfaces with different triangulations and sets of edge lengths may nonetheless be isometric—consider for instance splitting the square




FIGURE 5. A given polyhedron admits many different geodesic triangulations, all of which describe the same intrinsic geometry.




FIGURE 6. A vertex of a polyhedron (left) is intrinsically the same as a circular cone (right), and hence experiences infinite scale distortion u when conformally flattened in the classical sense. (Black lines depict pullback of integer grid lines under a conformal flattening.)

faces of a cube along different diagonals. From this point of view, the intrinsic geometry of a discrete surface is completely determined by the number of vertices, along with their locations and cone angles. This situation motivates the construction of *intrinsic Delaunay triangulations* (Setion 5.2), which play an important role in the theory of discrete uniformization (Setion 5.3).

3. Part I: Discretized Conformal Maps

In Part I we consider several candidate ways to discretize conformal maps $f : (M, g) \rightarrow \mathbb{C}$ from a disk-like domain (M, g) to the complex plane \mathbb{C} , each of which fails to capture the behavior of smooth conformal maps in some essential way. The first question to ask, perhaps, is: why do we need a new definition at all? After all, a “discrete” surface is still a Riemannian manifold, and we should therefore be able to apply standard definitions directly (perhaps with some extra care in order to deal with cone points). This viewpoint of conformally equivalent cone metrics has in fact been studied by Troyanov [Tro91], though for our purposes it is not the right philosophical starting point: it views a polyhedron as a literal description of the geometry we want to represent, rather than a proxy for a smooth surface. From the classical Riemannian point of view a convex polyhedron does not look much

like a smooth surface: it is flat almost everywhere, except at a discrete collection of points where it locally looks like a cone. As depicted in Figure 6, a conformal flattening of such a metric hence looks nothing like a flattening of a smooth surface, since it has unbounded scale distortion u in the vicinity of each vertex. We are therefore motivated to seek an alternative definition for discrete conformal maps, which more faithfully captures the behavior we expect from smooth surfaces.

The notions studied in this part all begin with a fairly naive hypothesis (*e.g.*, “discrete conformal maps should preserve angles”), which in each case leads to a definition that is significantly more rigid than the smooth definition. An important exception is the finite element treatment discussed in Section 3.5, which provides the expected degree of flexibility, but otherwise does not capture many of the basic features of smooth conformal maps in an exact sense—for instance, even just composing a given finite element solution with a Möbius transformation may yield a map that is not the solution to any finite element problem. In all other cases, one can still consider maps which come as close as possible to satisfying the naive definition, *e.g.*, in the “least-squares” sense—until fairly recently, this point of view has been the starting point for most practical computational algorithms [LPRM02, SLMB05].

3.1. Interior Angle Preservation. We first consider the most basic geometric characterization on conformal maps, namely that they should preserve angles. In the smooth setting, an immersion $f : M \rightarrow \mathbb{C}$ is conformal if at every point $p \in M$ the angle between any two tangent vectors $X, Y \in T_p M$ is the same as the angle between their images $df_p(X), df_p(Y)$ in the plane. A very tempting idea, then, is to seek a discrete immersion $f : M \rightarrow \mathbb{C}$ that preserves the interior angles θ_i^{jk} at every triangle corner. However, it quickly becomes clear that this notion of conformal equivalence is far too rigid: interior angles are preserved if and only if each triangle $ijk \in F$ experiences a similarity transformation, or (intrinsically) if the new and old edge lengths in each triangle ijk are related by a scale factor $\lambda_{ijk} > 0$, *i.e.*, $\tilde{\ell}_{ab} = \lambda_{ijk} \ell_{ab}$ for $ab \in ijk$. At first glance this idea of locally scaling the discrete metric in each triangle feels reminiscent of the scaling relationship $\tilde{g} = e^{2u} g$ between conformally equivalent smooth metrics g, \tilde{g} . However, since each interior edge ij is shared by two triangles ijk, jil , the scale factors $\lambda_{ijk}, \lambda_{jil}$ of any two adjacent triangles must be identical. Globally, then, any piecewise linear map that preserves all interior angles is necessarily an isometry, up to a uniform scaling by some constant factor $c \in \mathbb{R}_{>0}$. Hence, an angle-based definition of conformal equivalence is far too rigid: each equivalence class is a one-parameter family of discrete metrics, parameterized by a single constant c . In stark contrast, smooth equivalence classes are parameterized by a scalar *function*, namely the log conformal factor $u : M \rightarrow \mathbb{R}$. (Yet as we will see in Part III a different notion of metric scaling *will* ultimately lead to the right notion of discrete conformal equivalence.)

One can nonetheless consider the map that distorts angles the least, *i.e.*, the one that minimizes the deviation of angles from their original values in the least-squares sense [SS00]. To do so, suppose we parameterize piecewise linear maps $f : M \rightarrow \mathbb{C}$ via the interior angles $\tilde{\theta}_i^{jk} > 0$ of the image $f(M)$. These angles describe a valid planar triangulation as long as they satisfy a collection of discrete integrability conditions:

- (1) For each triangle $ijk \in F$, $\tilde{\theta}_i^{jk} + \tilde{\theta}_j^{ki} + \tilde{\theta}_k^{ij} = \pi$.
- (2) For each interior vertex $i \in V$, $\sum_{ijk \in \text{St}(i)} \tilde{\theta}_i^{jk} = 2\pi$.




FIGURE 7. Angles that exhibit Euclidean sums around triangles and vertices may still fail to characterize a valid triangulation.

- (3) For each interior vertex $i \in \mathcal{V}$, $\prod_{ijk \in \text{St}(i)} \frac{\sin \tilde{\theta}_j^{ki}}{\sin \tilde{\theta}_k^{ij}} = 1$.

The first condition is the usual Euclidean triangle postulate; the second is simply the requirement that every vertex have a Euclidean angle sum. As illustrated in Figure 7, the third condition is a necessary closure condition on edge lengths around each vertex, obtained by applying the law of sines to the relationship

$$\prod_{ij \in \text{St}(i)} \frac{\tilde{\ell}_{i,j}}{\tilde{\ell}_{i,j+1}} = 1,$$

where $\tilde{\ell}_{i,j}$ and $\tilde{\ell}_{i,j+1}$ denote the lengths of consecutive edges in $\text{St}(i)$ with respect to counter-clockwise ordering. The best approximation of an angle-preserving map (in the least-squares sense) is then any minimizer of the energy

$$E_{\text{ang}}(\tilde{\theta}) := \sum_{ijk \in \mathcal{F}} (\tilde{\theta}_i^{jk} - \theta_i^{jk})^2,$$

subject to the condition that the new angles $\tilde{\theta}$ are positive and satisfy the discrete integrability conditions outlined above. These angles determine a discrete metric (up to a global uniform scaling), which in turn determines a discrete map to the plane (up to Euclidean motions). Here again we observe too much rigidity in the sense that E_{ang} typically has only a single unique minimizer (up to Euclidean motions), whereas in the smooth setting there is a large family of conformal maps from any disk-like domain to the plane. However, this minimizer will at least be a discrete immersion (due to condition (2) and the positive angle condition), and will exhibit variable rather than uniform scaling. Further properties of such maps are not well-understood—empirically, for instance, the minimizer of E_{ang} appears to approximate the conformal map of least *area* distortion, though this observation has never been carefully analyzed. A variety of numerical algorithms for computing such maps have been developed [SS00, SLMB05, ZLS07].

3.2. Cauchy-Riemann Equation. Analytically, smooth conformal maps on a domain $U \subset \mathbb{C}$ are traditionally characterized by the *Cauchy-Riemann equation*, which says that a map $f : U \rightarrow \mathbb{C}; (x, y) \mapsto u(x, y) + w(x, y)$ is holomorphic if

$$(3.1) \quad \frac{\partial u}{\partial x} = \frac{\partial v}{\partial y} \quad \text{and} \quad \frac{\partial u}{\partial y} = -\frac{\partial v}{\partial x}.$$




FIGURE 8. Geometrically, the Cauchy-Riemann equation asks that pushing forward vectors via given map f commutes with scaling and rotation.

If f is also an immersion (*i.e.*, f has a nonvanishing derivative), then it is *conformal*. On regular quadrilateral nets, where one has clear “ x ” and “ y ” directions, this coordinate description is a natural starting point for structure-preserving discretizations of holomorphic maps that capture much of the rich structure found in complex analysis [BG17]. For general triangulated surfaces, where there are no clearly distinguished coordinate directions, one must take a different approach. An alternative way to write Equation (3.1) is

$$df(\iota X) = \iota df(X),$$

where ι denotes the imaginary unit and X is any tangent vector field on U . In other words, a complex map f is holomorphic if pushing forward vectors commutes with 90-degree rotation. For a surface M with linear complex structure \mathcal{J} , a map $f : M \rightarrow \mathbb{C}$ is likewise holomorphic if

$$(3.2) \quad df(\mathcal{J}X) = \iota df(X)$$

for all tangent vector fields X on M .

When applied to piecewise linear maps $f : M \rightarrow \mathbb{C}$, this characterization again leads to a definition for discrete conformal maps that is far too rigid. In particular, the only affine maps that satisfy Cauchy-Riemann on a single triangle are Euclidean motions and uniform scaling. We therefore encounter the exact same situation as we did for angle preservation in Section 3.1: in order for affine maps to agree across shared edges, they must all exhibit identical scaling. Hence, a piecewise linear map f is holomorphic in the sense of Equation (3.2) if and only if it is a global isometry, up to a global uniform dilation.

As with angle preservation, one can nonetheless seek the piecewise linear map $f : V \rightarrow \mathbb{C}$ that best agrees with the Cauchy-Riemann equation. Consider for instance the energy

$$(3.3) \quad E_C(f) := \int_M |*df - \iota df|^2 dA,$$

which measures the L^2 residual of Equation (3.2); here, $*$ denotes the Hodge star operator on differential 1-forms. To avoid the trivial solution $f = 0$, one typically incorporates a pair of point constraints $f(a) = 0, f(b) = 1$ (for $a, b \in M, a \neq b$). Minimizing this energy over piecewise linear maps f leads to a standard




FIGURE 9. Top: two critical points f_1, f_2 of discrete Dirichlet energy among piecewise linear maps f from a triangulated disk M to the unit circular disk in the plane. Bottom: maps visualized by pulling back coordinate lines under f_1, f_2 .

Galerkin finite element method, yielding a corresponding discrete energy $\hat{E}_C(f)$ (see [CdGDS13, Chapter 7] for a detailed derivation); this approach was one of the earliest starting points for practical conformal surface parameterization [LPRM02]. As with E_{ang} , the (constrained) discrete energy \hat{E}_C typically has a unique minimizer. In contrast, the smooth energy E_C has an enormous space of minimizers—consider composing any minimizer f with an in-plane conformal map $\eta : f(M) \rightarrow \mathbb{C}$. Hence, the unique minimizer chosen by \hat{E}_C will depend in an unstable way on superficial features of the problem (such as the choice of tessellation), hinting at some of the practical difficulties with definitions that are “too rigid”. For further discussion, see [MTAD08, Section 3] and [SC17, Section 2].

3.3. Critical Points of Dirichlet Energy. The *Dirichlet energy* of a differentiable map $f : M_1 \rightarrow M_2$ between Riemannian manifolds (M_1, g_1) and (M_2, g_2) is the functional

$$E_D(f) := \int_{M_1} |df|^2 dV,$$

where dV is the volume measure on M_1 ; f is *harmonic* if it is a critical point of E_D . Purely topological conditions determine whether a harmonic map is also holomorphic:

THEOREM 3.1 (Eells & Wood 1975). *If $f : M_1 \rightarrow M_2$ is a harmonic map and*

$$\chi(M_1) + |\deg(f)\chi(M_2)| > 0,$$

then f is either holomorphic or antiholomorphic (where $\chi(M)$ is the Euler characteristic of M , and $\deg(f)$ is the topological degree of f).

Antiholomorphic simply means angle preserving and orientation *reversing*, rather than orientation *preserving*. For instance, any harmonic map from the sphere to itself automatically preserves angles, since in this case $\chi(M_1) = \chi(M_2) = \chi(S^2) = 2$. Likewise, any harmonic map f from a topological disk M to the unit disk $D^2 \subset \mathbb{C}$ will be angle-preserving, since here $\chi(M) = \chi(D^2) = 1$. Hence, we can attempt

to define discrete holomorphic maps as those that minimize some kind of discrete Dirichlet energy.

For a discrete surface $M = (V, E, F)$ with metric $\ell : E \rightarrow \mathbb{R}_{>0}$, the discrete Dirichlet energy of a map $f : V \rightarrow \mathbb{C}$ can be defined as

$$(3.4) \quad \widehat{E}_D(f) := \sum_{ij \in E} w_{ij} |f_j - f_i|^2$$

for any collection of edge weights w_{ij} which make this energy a positive-semidefinite quadratic form (*i.e.*, $\widehat{E}_D(f) \geq 0$ for all piecewise linear maps f). Letting k and l be the vertices opposite edge ij , a common choice is the *cotangent weights* $w_{ij} := \frac{1}{2}(\cot \theta_k^{ij} + \cot \theta_l^{ji})$, which arise from the finite element discretization mentioned in Section 3.2. (Other choices are also possible—see for instance [HS15].) A discrete holomorphic map to a fixed region $U \subset \mathbb{C}$ can then be defined as a critical point of \widehat{E}_D , subject to the condition that boundary vertices are mapped to ∂U ; this map will look conformal (rather than just holomorphic) if the boundary polygon has a unit turning number—see Figure 9 for two examples. This point of view was originally considered by Hutchinson [Hut91]; see also [PP93].

A closely related point of view is that, in the smooth setting, the conformal energy $E_C(f)$ of any map $f : M \rightarrow \mathbb{C}$ can be expressed as the difference between the Dirichlet energy $E_D(f)$ and the signed area $\mathcal{A}(f)$ of the image $f(M)$, *i.e.*,

$$E_C(f) = E_D(f) - \mathcal{A}(f)$$

(see [CdGDS13, Chapter 7]). In the case where the target is fixed (*e.g.*, if one considers only maps to the unit disk), the area term \mathcal{A} is constant, and hence E_D will have the same minimizers as E_C . One can hence define a discrete conformal map f as a minimizer of the discrete Dirichlet energy \widehat{E}_D (Equation (3.4)) minus the signed area of the target polygon (subject to the same point constraints as in Section 3.2), giving another quadratic form in f :

$$\widehat{\mathcal{A}}(f) := \frac{1}{2} \sum_{ij \in \partial M} f_i \times f_j;$$

here ∂M denotes the collection of oriented edges in the boundary of M , and $z_1 \times z_2 := \text{Im}(\bar{z}_1 z_2)$. This formulation is used as the starting point for several practical algorithms in digital geometry processing [DMA02, MTAD08]. It turns out, however, that the energy $\widehat{E}_D - \widehat{\mathcal{A}}$ is identical to the energy \widehat{E}_C obtained by discretizing the Cauchy-Riemann equations [CSD02]. Its minimizers therefore exhibit the same degree of rigidity as before.

3.4. Hodge Duality. Consider the Hodge star $*$ on differential 1-forms α , which on a smooth surface with linear complex structure \mathcal{J} can be expressed via the relationship¹

$$*\alpha(X) = \alpha(\mathcal{J}X)$$

for all tangent vector fields X . Since conformal maps preserve the linear complex structure, they also preserve the 1-form Hodge star. In the discrete setting, one can reverse this relationship and try to *define* a discrete conformal map as one that preserves the (discrete) Hodge star [Mer01]. In particular, differential forms can be discretized as cochains on a simplicial manifold and its polyhedral dual [Wil05, DHLM05]. A fairly common discretization of the Hodge star is then a “diagonal”

¹Note that some authors adopt an opposite orientation convention, *i.e.*, $*\alpha(X) = -\alpha(\mathcal{J}X)$.




FIGURE 10. The ratio of dual over primal edge length is equal to half the sum of the cotangents of the opposite angles θ_k^{ij} , θ_l^{ji} .

linear map from primal k -cochains to dual $(n - k)$ -cochains determined by the ratio of primal and dual volumes. In the particular case of discrete 1-forms on a triangulated surface and its circumcentric dual, a discrete differential 1-form can be encoded via a value α_{ij} per primal edge, and the corresponding (Hodge) dual 1-form can be expressed as a value on each dual edge via the cotangent formula:

$$(3.5) \quad * \alpha_{ij} := \underbrace{\frac{1}{2}(\cot \theta_k^{ij} + \cot \theta_l^{ji})}_{w_{ij}} \alpha_{ij},$$

where θ_k^{ij} , θ_l^{ji} are the angles opposite ij , as depicted in Figure 10. A discrete conformal map in the sense of Hodge duality is then any piecewise linear map for which the new angles $\tilde{\theta}$ satisfy $\cot \tilde{\theta}_k^{ij} + \cot \tilde{\theta}_l^{ji} = 2w_{ij}$, or, from an intrinsic point of view, any new assignment of edge lengths $\tilde{\ell}_{ij}$ that preserves this same quantity.

Initially, one might be optimistic that a discretization based on the Hodge star yields less rigidity than one based on preservation of angles, since (at least naively) preservation of the edge weights w_{ij} effectively places only $|\mathbf{E}| \approx 3|\mathbf{V}|$ conditions on the map f , whereas preservation of interior angles θ_i^{jk} corresponds to $3|\mathbf{F}| \approx 6|\mathbf{V}|$ constraints. However, recent analysis [ZGLG12, GMMD14] extinguishes any such optimism:

THEOREM 3.2. *The primal-dual length ratios $w_{ij} := \frac{1}{2}(\cot \alpha_{ij} + \cot \beta_{ij})$ on a given discrete surface $\mathbf{M} = (\mathbf{V}, \mathbf{E}, \mathbf{F})$ uniquely determine the discrete metric $\ell : \mathbf{E} \rightarrow \mathbb{R}$, up to global scaling.*

The proof relies on the fact that a metric exhibiting a prescribed Hodge star can be obtained as a minimizer of a convex function, where the Hessian at any such minimizer has a kernel consisting only of vectors corresponding to global scaling of ℓ —exactly the same degree of rigidity as with angle preservation (Setion 3.1), Cauchy-Riemann (Setion 3.2), and Dirichlet energy (Setion 3.3).

3.5. Conjugate Harmonic Functions. Suppose we express a holomorphic map $f : M \rightarrow \mathbb{C}$ as $f = a + bi$ for a pair of real-valued functions $a, b : M \rightarrow \mathbb{R}$. A straightforward consequence of the Cauchy-Riemann equation (Setion 3.2) is that a and b form a *conjugate harmonic pair*, i.e., they are real harmonic functions with orthogonal gradients:

$$(3.6) \quad \begin{aligned} \Delta a &= 0, \\ \Delta b &= 0, \\ \nabla b &= \mathcal{J} \nabla a, \end{aligned}$$




FIGURE 11. One possible notion of discrete conjugacy: a function on primal (black) vertices is conjugate to a function on dual (white) vertices if the difference across primal edges is equal to the difference across dual edges, up to a constant factor that accounts for triangle shape.

Here Δ is the Laplace-Beltrami operator on M , ∇ is the gradient operator, and \mathcal{J} is the linear complex structure.

The conditions in Equation (3.6) are of course closely related to (in fact, equivalent to) the condition that the map itself be a critical point of Dirichlet energy (Section 3.3), but the perspective of conjugate harmonic pairs will provide a different starting point for discretization, where boundary conditions play a key role. Characterizing conformal maps in terms of harmonic functions is also attractive from the perspective of discretization, since discrete harmonic functions are well-studied. In the simplicial setting, a discrete harmonic function $\phi : V \rightarrow \mathbb{R}$ is naturally defined as a piecewise linear function in the kernel of a discrete Laplace-Beltrami operator $L : V \rightarrow V$ such as the *cotan Laplacian*

$$(3.7) \quad (L\phi)_i := \frac{1}{2} \sum_{ij \in E} (\cot \theta_k^{ij} + \cot \theta_l^{ji})(\phi_j - \phi_i).$$

(See [WMKG07] for a more thorough discussion.) To define a conjugate harmonic pair, one then just needs a discrete notion of conjugacy.

What does it mean for two discrete harmonic functions to be conjugate? One idea, studied by Mercat and others [Mer01], is to consider functions on the combinatorial or *Poincaré dual* of a discrete surface $M = (V, E, F)$, which associates each vertex with a 2-cell, each edge with a 1-cell, and each triangle with a 0-cell. In this setting, two real-valued functions \mathbf{a}, \mathbf{b} on the primal and dual 0-cells (*resp.*) are discretely conjugate if for each edge e and corresponding dual edge e^* , the difference of \mathbf{a} values across e is equal to the difference of \mathbf{b} values across e^* , up to a scale factor $w_e \in \mathbb{R}$ that accounts for the geometry of the triangulation (as discussed in Section 3.4). In other words, if

$$\mathbf{b}_{j^*} - \mathbf{b}_{i^*} = w_{ij}(\mathbf{a}_j - \mathbf{a}_i)$$


for all edges $ij \in E$, where i^*, j^* are the associated dual vertices (see Figure 11). In terms of (discrete) differential forms, we are simply requiring that $db = *da$, capturing the conjugacy condition $\nabla b = \mathcal{J}\nabla a$. Here one encounters two basic sources of difficulty. First, as detailed in Section 3.4, if one views the conformal structure of a discrete surface as being determined by the edge weights w_{ij} (or *discrete Hodge star*), then one encounters severe rigidity: the weights uniquely determine a discrete metric. Moreover, pairs of functions \mathbf{a}, \mathbf{b} that are conjugate in

this sense do not define a *simplicial* mapping from M to \mathbb{C} , since the coordinates \mathbf{b} are associated with 0-cells of the dual complex, rather than vertices of the original discrete surface M .

Alternatively, one can take the following approach: given a discrete harmonic function \mathbf{a} (*i.e.*, a piecewise linear function satisfying $L\mathbf{a} = 0$), its harmonic conjugate can be defined as the function $\mathbf{b} : V \rightarrow \mathbb{R}$ that minimizes the L^2 difference between $\mathcal{J}\nabla\mathbf{a}$ and $\nabla\mathbf{b}$ over the space of piecewise linear functions—or equivalently, which minimizes the discrete conformal energy $\widehat{E}_C(\mathbf{f}) = \widehat{E}_D(\mathbf{f}) - \widehat{A}(\mathbf{f})$ of a map $\mathbf{f} = \mathbf{a} + \mathbf{b}i$ (Setion 3.3), while keeping \mathbf{a} fixed. The minimizer is the solution to the discrete Laplace equation $L\mathbf{b} = 0$ with discrete Neumann boundary data

$$h_i = \frac{1}{2}(a_{i+1} - a_{i-1}),$$

where $i-1$, i , and $i+1$ denote consecutive vertices along the boundary (see [SC17, Section 4.3.3] for more detail). The resulting map $\mathbf{f} := \mathbf{a} + \mathbf{b}i$ is then discrete conformal in the same finite element sense discussed in Setion 3.2. Importantly, however, one finally obtains exactly the right amount of flexibility: whereas a “least squares” conformal map [LPRM02] of the kind discussed in Setion 3.2 is determined up to Euclidean motions, one now obtains a whole family of discrete conformal maps \mathbf{f} parameterized by the harmonic function \mathbf{a} . Since this function is in turn determined by its own boundary values, one ends up in precisely the same situation as in the smooth setting, where holomorphic maps $f : M \rightarrow \mathbb{C}$ can be parameterized by real functions on ∂M specifying either Dirichlet or Neumann boundary conditions. A more geometric view is that one can parameterize such maps by specifying either the log conformal factor u or the geodesic curvature κ along the boundary—Sawhney & Crane [SC17] outline a complete strategy in the piecewise linear setting.



Ultimately, it should come as no surprise that a proper finite element treatment should faithfully capture the behavior of smooth conformal maps *in the limit of refinement*. On the other hand, for any *fixed* triangulation, such maps do not exactly preserve many basic properties from the smooth setting, such as covariance with respect to Möbius transformations. Moreover, they do not provide a natural notion of equivalence—for instance, the composition of two piecewise linear harmonic functions is not in general harmonic. We will therefore continue in Parts II and III to seek a precise notion of discrete conformal equivalence for finite triangulations.

4. Part II: Circle Preservation

A linear map preserves angles if and only if it is the composition of a rotation and a dilation (no shear); hence, it also preserves circles. Since at each point $p \in M$ the differential df_p of a conformal map f is an angle-preserving linear map, it must also preserve infinitesimal circles. This point of view is the starting point for several distinct but closely-related approaches to discrete conformal maps, based on arrangements of circles with special combinatorial and geometric relationships.

Finite arrangements of circles provide fertile soil for discrete conformal maps, capturing many features of the smooth theory. Koebe originally showed how incidence relationships in planar graphs can be captured via tangency relationships between circles, providing a basic connection between combinatorics and geometry.




FIGURE 12. Left: a smooth Riemann mapping from a simply connected region in the plane to the unit circular disk. Right: a discrete Riemann mapping, expressed as a *circle packing*.

Thurston conjectured (and Rodin & Sullivan later proved) that these *circle packings* have a deep connection to Riemann mappings in the complex plane (Section 4.1); such connections have since been studied intensively by Schramm [Roh11] and many others. To generalize this construction to curved surfaces and irregular triangulations one must incorporate additional geometric information, such as intersection angles or other data describing relationships between circles; such arrangements fall under the general heading of *circle patterns*² (Section 4.2).

There are however some missing features, most notably the lack of a uniformization theorem that guarantees the existence of a constant curvature circle pattern equivalent to any given triangulation. In Part III we will see a discrete uniformization theorem which guarantees existence by considering variable rather than fixed triangulations; whether one can take an analogous approach in the circle pattern setting remains to be seen. Note that circle packings and patterns have been studied extensively beyond the context of triangulations and discrete conformal maps; see for instance [Ste05, BS04a, BHS06, WP11].

4.1. Circle Packings. A *circle packing* is a collection of closed circular disks in the plane (or other surface) that intersect only at points of tangency—see [Ste03] for an excellent overview. To any such collection one can associate a graph $G = (V, E)$ (called the *nerve*) where each vertex corresponds to a disk, and two vertices are connected by an edge if and only if their associated disks are tangent. A natural question to ask is: which graphs admit circle packings?

THEOREM 4.1 (Circle Packing). *Every planar graph $G = (V, E)$ can be realized as a circle packing in the plane.*

(For nonplanar graphs, one can also consider circle packings on surfaces of higher genus.) A first hint that circle packings are connected to discrete conformal maps comes via the following theorem:

THEOREM 4.2 (Koebe). *If G is a connected maximal planar graph, then it has a unique circle packing up to Möbius transformations and reflections.*

(A constructive algorithm for computing such packings was given by Collins & Stephenson [CS03].) A graph is *maximal* if the addition of any edge makes it

²There is some inconsistency in the use of *packings* versus *patterns* throughout the literature—for clarity we adopt the convention that *packings* are arrangements of circles meeting tangentially, whereas *patterns* describe any arrangement of circles (possibly overlapping or disjoint), whether associated with vertices, faces, etc.




FIGURE 13. Discrete surfaces with identical combinatorics but different edge lengths (left) yield the same class of circle packings into the unit disk (right).

nonplanar. If G is also finite, one quickly sees that it is equivalent to a triangulation of the sphere, via stereographic projection. Koebe's theorem therefore says that any triangulated sphere can be realized as a family of planar circle packings parameterized by Möbius transformations, just as smooth conformal maps from the sphere to the plane have Möbius symmetry. An even richer connection between circle packings and conformal maps can be made via the *Riemann mapping theorem*:


THEOREM 4.3 (Riemann mapping). *Any nonempty simply connected open set $\Omega \subsetneq \mathbb{C}$ can be mapped to the interior of the unit circular disk $D^2 := \{z \in \mathbb{C} : |z| < 1\}$ by a bijective map $\phi : \Omega \rightarrow D^2$ that is holomorphic (hence conformal) in both directions.*

The circle packing analogue of Riemann mapping was originally conjectured by Thurston, and later proved by Rodin and Sullivan [RS87]. The basic idea is to start with a regular hexagonal circle packing C of a simply connected region Ω by disks of radius $\varepsilon > 0$, *i.e.*, for any hexagonal tiling of the plane, take only those disks that intersect Ω . Now find a circle packing C' that maintains the same incidence relationships, but where all disks along the boundary are now tangent to the unit circular disk D^2 (this idea is illustrated in Figure 12, right). The relationship between these two packings defines an *approximate mapping* of Ω to D^2 : for sufficiently small ε any point $z \in \Omega$ will be contained in a circle c from C , and can be mapped to the center of the corresponding circle c' from C' . Rodin and Sullivan show that the approximate mapping converges to a conformal homeomorphism as $\varepsilon \rightarrow 0$.

Unlike smooth Riemann mappings, however, one cannot directly use circle packings to define discrete conformal maps between any two disk-like regions, since finite hexagonal packings of these regions will not in general have the same combinatorics. More importantly, this theory does not provide a general approach to uniformization, since it cannot account for the curvature of the domain, nor triangulations with nonuniform edge lengths. Consider for instance simplicial disks with identical combinatorics but different discrete metrics $\ell : E \rightarrow \mathbb{R}_{>0}$, as depicted in Figure 13. Since a circle packing depends purely on the combinatorics of the edge graph, these surfaces are realized by identical families of circle packings—implying

that all discrete metrics on a given simplicial disk are, effectively, conformally equivalent. In this sense, the basic theory of circle packings is *too flexible*; one is therefore prompted to enrich it with additional data.

4.2. Circle Patterns. A more direct way to encode the geometry of the domain is to incorporate additional metric information into the arrangement of circles. Suppose we associate each edge $ij \in E$ of a triangulation $M = (V, E, F)$ with an angle $\omega_{ij} \in [0, \frac{\pi}{2}]$; the pair (M, ω) is called a *weighted triangulation*. If we now assign a radius $r_i > 0$ to each vertex $i \in V$, then a pair of circles in the Euclidean plane with the prescribed intersection angle ω_{ij} and radii r_i, r_j will have centers separated by a distance



$$(4.1) \quad \ell_{ij} = \sqrt{r_i^2 + r_j^2 + 2r_i r_j \cos \omega_{ij}}.$$

The condition $\omega \leq \pi/2$ ensures that these lengths satisfy the triangle inequality in each face $ijk \in F$ (see [Thu79, Lemma 13.7.2]); they hence determine a discrete metric on the surface (called the *circle packing metric*), and in turn, a corresponding cone angle Ω_i at each vertex (Equation (2.1)). We can make an analogy between each of these quantities and data from the smooth setting:

- the edge lengths ℓ_{ij} encode the metric,
- the cone angles Ω_i encode the Gaussian curvature,
- the intersection angles ω_{ij} capture the conformal structure, and
- the radii r_i (or more precisely, any subsequent *changes* to the radii) play the role of conformal scale factors.

The last two items on this list can be understood by considering that (i) angles are conformal invariants, and (ii) adjusting the radii (scale factors) will scale the edge lengths ℓ (metric) while keeping the angles ω (conformal structure) fixed. From this point of view, finding a conformal map from a curved surface ($\Omega \neq 0$) to a flat one ($\Omega_i = 0$ at all vertices) amounts to finding an appropriate adjustment of radii, *i.e.*, finding the conformal scale factors that flatten the metric. Note however that it is not immediately clear that a target angle defect Ω can always be achieved—and in general it cannot. In particular, we have the following theorems [Thu79]:

THEOREM 4.4 (rigidity of circle patterns). *For a closed weighted triangulation (M, ω) , an assignment of cone angles Ω_i to vertices uniquely determines a unique discrete metric ℓ_{ij} if one exists—in other words, if there are radii r_i at vertices such that (i) circles intersect at the prescribed angles ω and (ii) the edge lengths determined by Equation (4.1) induce the prescribed cone angles Ω , then this circle pattern is unique up to a uniform scaling of all circle radii / edge lengths.*


THEOREM 4.5 (existence of circle patterns). *Let $M = (V, E, F)$ be a closed weighted triangulation with weights ω , and for any subset of vertices $\mathfrak{l} \subset V$ let $\text{Lk}(\mathfrak{l})$ and $F_{\mathfrak{l}}$ denote the simplicial link of \mathfrak{l} and the set of faces with vertices in \mathfrak{l} , resp. Then prescribed cone angles Ω_i can be achieved if and only if*

$$\sum_{i \in \mathfrak{l}} \Omega_i + \sum_{ij \in \text{Lk}(\mathfrak{l})} (\pi - \omega_{ij}) > 2\pi\chi(F_{\mathfrak{l}}),$$

where χ denotes the Euler characteristic.

In short: one cannot prescribe curvature arbitrarily, but if the prescribed curvature can be achieved, then it uniquely determines the metric (up to global scaling). The condition in Theorem 4.5 is akin to Gauss-Bonnet, but much stronger: just as any metric can be uniformized in the smooth setting, one would like the existence of discrete uniformization to depend only on a basic Gauss-Bonnet condition on the curvature data Ω , and not on the data ω describing the domain itself.

Generalizations. Generalizations of the circle patterns described above provide greater flexibility. One such generalization, first studied by Bowers and Stephenson [BS04b], is an *inversive distance circle packing*—here we exchange the angles ω_{ij} for values $I_{ij} \in [-1, \infty)$, and define the edge lengths via the *inversive distance*




$$\ell_{ij} = \sqrt{r_i^2 + r_j^2 + 2r_i r_j I_{ij}}.$$

When $I_{ij} = 1$ circles are tangent, when $I_{ij} < 1$ they intersect at an angle $\omega_{ij} = \arccos(I_{ij})$, and when $I_{ij} > 1$ they are disjoint. In this setting one again has local and global rigidity theorems akin to Theorem 4.4 [Guo11, Luo11, Xu18], though no guarantee of existence due to conditions like the one from Theorem 4.5.

One can also associate a circle pattern with the *faces* of a triangulation, rather than its vertices. Let (M, ω) be a weighted triangulation where (i) the weights ω_{ij} sum to 2π around each interior vertex, and (ii) the values $\kappa_i := 2\pi - \sum_{ij} \omega_{ij}$ sum to 2π over all boundary vertices. One can then find a Euclidean circle pattern with one circle per face, so long as there is a collection of interior angles α_i^{jk} that satisfy Rivin’s *coherent angle system* [Riv94], namely

- (1) (Positivity) $\alpha_i^{jk} > 0$ at each triangle corner.
- (2) (Triangle Postulate) $\alpha_i^{jk} + \alpha_j^{ki} + \alpha_k^{ij} = \pi$ for each face ijk .
- (3) (Compatibility) $\pi - \omega_{ij}$ is equal to $\alpha_k^{ij} + \alpha_l^{ji}$ for each interior edge ij , and equal to α_k^{ij} for each boundary edge ij .

For instance, if ω_{ij} are the intersection angles of the triangle circumcircles in a planar Delaunay triangulation, then the circle pattern will reproduce this triangulation, up to a global similarity. Existence and uniqueness (up to similarity) can be obtained through variational arguments [BS04a]. However, this construction does not directly address the task of uniformization, since for a curved domain the intersection angles ω_{ij} will not in general sum to 2π around interior vertices. A practical approach to conformal flattening via facewise circle patterns was proposed by Kharevych *et al.* [KSS06], who first construct a collection of interior angles α that (i) satisfy Rivin’s conditions, and (ii) approximate the interior angles θ_i^{jk} of the curved domain in the least-squares sense. This particular treatment does not however provide a clean notion of discrete conformal equivalence for general (*i.e.*, non-flat) discrete surfaces.



Other generalizations, such as *hyper-ideal circle patterns* [Sch08, Spr08], finally provide the right degree of flexibility. In this setting, one associates circles with both vertices *and* edges—and obtains a discrete uniformization theorem in the case of surfaces with non-positive curvature [BDS17], even with fixed combinatorics. Such results suggest that it is not unreasonable to expect a complete discrete uniformization theorem based on patterns of circles (namely, one that includes the positively curved case), though at present it would seem that perhaps the ending has not yet been written.

4.3. Discrete Ricci Flow. In the smooth setting, one approach to uniformization is to consider Hamilton’s *Ricci flow* [Ham82], which for surfaces can be expressed as

$$\frac{d}{dt}g = -Kg.$$

Here K denotes the Gaussian curvature of the current metric g . Intuitively, Ricci flow shrinks the metric where curvature is “too positive,” and expands it where curvature is “too negative,” ultimately smoothing out all bumps in curvature. Since the flow applies a pointwise rescaling at each moment in time, the final, uniformized metric is conformally equivalent to the initial one. Often one considers a normalized version of this flow

$$\frac{d}{dt}g = (\bar{K} - K)g$$

where \bar{K} is the average Gaussian curvature over the domain; this flow becomes stationary when the metric has constant curvature.

Chow and Luo [CL03] consider a discrete analogue of this flow based on weighted triangulations (with intersection angles $\omega_{ij} \in [0, \frac{\pi}{2}]$)—in particular, they define the (*normalized*) *combinatorial Ricci flow*

$$\frac{d}{dt}r_i = (\bar{\Omega} - \Omega_i)r_i,$$

where the values Ω_i are the cone angles of the circle packing metric determined by r_i , and $\bar{\Omega}$ is the average cone angle. If one considers the effect of the radii r on the discrete metric ℓ , this flow has the same basic behavior as Ricci flow: the metric is rescaled in order to smooth out bumps in curvature. Chow and Luo show that this flow is defined for all time and converges to a constant curvature metric *if one exists, i.e.*, the flow does not uniformize all initial metrics—the condition on existence is the same as the one given in Theorem 4.5. Nonetheless, the flow behaves well enough in practice to provide a starting point for a wide variety of algorithms in geometric computing [GY08]. Circle patterns can also be connected with other geometric flows such as Calabi flow [Ge18, ZLH⁺18], though again Thurston’s condition (Theorem 4.5) must of course hold in order to guarantee existence of a uniformized circle packing metric.

Luo also considered a closely related *combinatorial Yamabe flow* [Luo04], defined in terms of a different discrete analogue of conformal maps (*i.e.*, one that is not based on circle packings or patterns). This flow provides a starting point for the notion of discrete conformal equivalence that we will study in Part III, and ultimately, to a complete discrete uniformization theorem where existence can be guaranteed.

5. Part III: Metric Scaling

In the smooth setting, two Riemannian metrics g, \tilde{g} on a surface M are *conformally equivalent* if they are related by a positive scaling, *i.e.*, if $\tilde{g} = e^{2u}g$ for some function $u : M \rightarrow \mathbb{R}$ called the *log conformal factor*. Hence, e^u gives the length scaling, and e^{2u} is the area scaling. A *conformal structure* on M is then an equivalence class of metrics.

This section examines a definition for conformal equivalence of discrete metrics that resembles the one found in the smooth setting, and which ultimately allows one to formulate a discrete uniformization theorem for polyhedral surfaces (Setion 5.3) that mirrors the one for smooth Riemann surfaces. From an applied point of view discrete uniformization provides a principled starting point for tasks like regular surface remeshing [CZ17] and constructing maps between polyhedral surfaces [HK15]. We begin with a basic picture involving only Euclidean geometry (Setion 5.1), and will revisit this story through the lens of hyperbolic geometry in Part IV.

5.1. Conformally Equivalent Discrete Metrics. Recall that a *discrete metric* is an assignment of edge lengths $\ell : E \rightarrow \mathbb{R}_{>0}$ that satisfy the triangle inequality in each face, *i.e.*, $\ell_{ij} + \ell_{jk} \geq \ell_{ki}$ for all $ijk \in F$. What does it mean for two discrete metrics $\ell, \tilde{\ell}$ to be conformally equivalent? One idea is to simply mimic the smooth definition and ask that $\tilde{\ell}_{ij} = \lambda_{ij}\ell_{ij}$ for some scale factor λ_{ij} on each edge $ij \in E$. However, this constraint is clearly “too flexible,” since then every pair of discrete metrics is conformally equivalent—simply let λ_{ij} be the ratio $\tilde{\ell}_{ij}/\ell_{ij}$. Also recall from Setion 3.1 that a scale factor *per face* is “too rigid,” since one is then forced to make all scale factors equal. An alternative is to consider scale factors at vertices [RW84, Luo04]:

DEFINITION 5.1. Two discrete metrics $\ell, \tilde{\ell}$ on the same triangulation $M = (V, E, F)$ are *discretely conformally equivalent* if for each edge $ij \in E$,

$$(5.1) \quad \tilde{\ell}_{ij} = e^{(u_i+u_j)/2}\ell_{ij},$$

for some collection of values $u : V \rightarrow \mathbb{R}$.

This definition again gives the impression of merely aping the smooth relationship—yet in this case the resulting theory is neither too rigid nor too flexible. Instead, it provides a rich notion of discrete conformal equivalence that beautifully preserves much of the structure found in the smooth setting [BPS15, GLSW18], a perspective which has had a significant impact on (and has been inspired by) algorithms from digital geometry processing [SSP08, BCGB08, CZ17].

One elementary observation is that, *locally*, conformally equivalent discrete metrics in the sense of Equation (5.1) are still very flexible:

LEMMA 5.2. *Any two discrete metrics $\ell, \tilde{\ell}$ on a single triangle ijk are discretely conformally equivalent.*

PROOF. The metrics are conformally equivalent if there exists an assignment of log scale factors u_i, u_j, u_k to the three vertices such that

$$e^{(u_a+u_b)/2} = \tilde{\ell}_{ab}/\ell_{ab}$$

for each edge $ab \in ijk$. Let $\lambda_{ij} := 2 \log(\ell_{ij})$ (and similarly for $\tilde{\ell}$). Then by taking the logarithm of the system above we obtain a linear system for the u values, namely

$$u_a + u_b = \tilde{\lambda}_{ab} - \lambda_{ab}, \quad \forall ab \in ijk.$$




FIGURE 14. A piecewise linear map is discrete conformal if and only if it preserves the length cross ratio $\mathbf{c}_{ij} := \ell_{il}\ell_{jk}/\ell_{ki}\ell_{lj}$ for each interior edge ij .

This system has a unique solution, independent of the values of ℓ and $\tilde{\ell}$. In particular,

$$(5.2) \quad e^{u_i} = \frac{\tilde{\ell}_{ij} \ell_{jk} \tilde{\ell}_{ki}}{\ell_{ij} \tilde{\ell}_{jk} \ell_{ki}},$$

and similarly for u_j, u_k . \square

At this point it might be tempting to believe that all discrete metrics are conformally equivalent, since they are equivalent on individual triangles. However, since scale factors (at vertices) are shared by several triangles, one still obtains an appropriate degree of rigidity. In particular, an important observation is that all metrics within the same discrete conformal equivalence class can be identified with a canonical collection of *length cross ratios* [SSP08]:

DEFINITION. Let $\ell : \mathbf{E} \rightarrow \mathbb{R}_{>0}$ be a discrete metric on a triangulated surface $\mathbf{M} = (\mathbf{V}, \mathbf{E}, \mathbf{F})$. For any pair of triangles $ijk, jil \in \mathbf{F}$ sharing a common edge $ij \in \mathbf{E}$, the associated *length cross ratio* is the quantity

$$(5.3) \quad \mathbf{c}_{ij} := \frac{\ell_{il}\ell_{jk}}{\ell_{ki}\ell_{lj}}.$$

THEOREM 5.3. *Two discrete metrics $\ell, \tilde{\ell}$ on the same triangulation \mathbf{M} are discretely conformally equivalent if and only if they induce the same length cross ratios $\mathbf{c}, \tilde{\mathbf{c}}$.*

PROOF. First suppose that $\ell, \tilde{\ell}$ are discretely conformally equivalent, *i.e.*, that $\tilde{\ell}_{ij} = e^{(u_i+u_j)/2}\ell_{ij}$ for some collection of log conformal factors $u_i : \mathbf{V} \rightarrow \mathbb{R}$. Then

$$\tilde{\mathbf{c}}_{ijkl} = \frac{\tilde{\ell}_{il}\tilde{\ell}_{jk}}{\tilde{\ell}_{ki}\tilde{\ell}_{lj}} = \frac{e^{(u_i+u_l)/2}e^{(u_j+u_k)/2}\ell_{il}\ell_{jk}}{e^{(u_k+u_i)/2}e^{(u_l+u_j)/2}\ell_{ki}\ell_{lj}} = \frac{\ell_{il}\ell_{jk}}{\ell_{ki}\ell_{lj}} = \mathbf{c}_{ijkl}.$$

Now suppose that $\ell, \tilde{\ell}$ induce identical cross ratios, *i.e.*, $\mathbf{c} = \tilde{\mathbf{c}}$. By Lemma 5.2, these metrics already satisfy the conformal equivalence relation on individual triangles. In particular, for any pair of adjacent triangles ijk, jil , we can find compatible log scale factors u_i, u_j, u_k and v_j, v_i, v_l . Moreover, these scale factors will agree on the shared edge ij if and only if $u_i = v_i$ and $u_j = v_j$. Applying Equation (5.2) we discover that this compatibility condition is equivalent to

$$\frac{\ell_{jk} \tilde{\ell}_{ki}}{\tilde{\ell}_{jk} \ell_{ki}} = \frac{\tilde{\ell}_{il} \ell_{lj}}{\ell_{il} \tilde{\ell}_{lj}},$$

i.e., equality of cross ratios. \square

The length cross ratios $\mathbf{c} : E \rightarrow \mathbb{R}$ effectively specify a point in the Teichmüller space of discrete metrics; see [BPS15, Remark 2.1.2] for further discussion. Extrinsicly, it is not hard to show that for a discrete surface immersed in \mathbb{R}^n , length cross ratios will be preserved by Möbius transformations of the vertices (assuming that the transformed vertices are connected by straight segments rather than circular arcs). As in the smooth setting, however, the space of maps $f : V \rightarrow \mathbb{C}$ that induce a discretely conformally equivalent metric $\tilde{\ell}_{ij} := |f_j - f_i|$ is much larger than just the space of Möbius transformations. One characterization is that discrete conformal maps (in this sense) correspond to piecewise projective maps that preserve triangle circumcircles [SSP08, Section 3.4]. Connections between conformally equivalent discrete metrics and piecewise Möbius transformations have also been studied [VMW15].

Two elementary examples help reinforce the connection between smooth and discrete conformal equivalence:

EXAMPLE 5.4 (discrete metrics on the two-triangle sphere). In the smooth setting, all 2-spheres have the same conformal structure. Consider the triangulation of the 2-sphere by two triangles depicted in Figure 2, bottom. Any discrete metric on this triangulation is determined by three edge lengths $\ell_{ij}, \ell_{jk}, \ell_{ki} > 0$ satisfying the triangle inequality. By the same argument as in Lemma 5.2, all discrete metrics on this triangulation of the sphere are discretely conformally equivalent—mirroring the situation in the smooth setting.

EXAMPLE 5.5 (discrete metrics on the two-triangle torus). In the smooth setting, there is a two-parameter family of conformal structures on the torus. Consider the triangulation of the torus depicted in Figure 2, top. Any discrete metric on this triangulation is determined by three edge lengths a, b, c corresponding to the edge marked with a single arrow, the edge marked with a double arrow, and the diagonal. Since there is only one vertex in the triangulation (and hence one scale factor), two such metrics are discretely conformally equivalent if and only if they are related by a uniform scaling—partitioning the three-parameter family of discrete metrics into a two-parameter family of conformal classes (as in the smooth setting).

There are however major difficulties with the theory discussed so far: it only allows us to discuss conformal equivalence of polyhedra with identical combinatorics, since it depends on comparing values associated with edges—at this point, for instance, it is not even possible to say that two different triangulations of the cube (with square faces split along different diagonals) are discretely conformally equivalent, even though they are isometric in the usual sense. Moreover, procedures for uniformizing a given discrete metric (discussed in Section 5.3) may break down before achieving constant curvature, *i.e.*, the metric may degenerate into a collection of edge lengths where the triangle inequality no longer holds. In order to develop a full theory of discrete uniformization, we therefore have to understand what it means for two combinatorially *inequivalent* polyhedra to be conformally equivalent. For this, we first need to consider a structure called the *intrinsic Delaunay triangulation*, which enables us to associate a canonical triangulation with any polyhedral cone metric.

5.2. Intrinsic Delaunay Triangulations. *Delaunay triangulations* arise naturally throughout discrete and computational geometry, and have numerous characterizations. For a triangulation in the plane, the most typical characterization is that

the circumcenter of every triangle is “empty,” *i.e.*, no vertices of the triangulation are contained in its interior. An equivalent condition is that the sum of angles opposite every interior edge must be no greater than π , *i.e.*, for any two triangles ijk, jil sharing an edge jil ,

$$(5.4) \quad \theta_k^{ij} + \theta_l^{ji} \leq \pi.$$

Unlike the empty circumcircle definition, the angle sum condition generalizes in a straightforward way to any triangulated surface $M = (V, E, F)$ with discrete metric $\ell : E \rightarrow \mathbb{R}_{>0}$, since interior angles θ are determined (via the law of sines) by the edge lengths ℓ . Any discrete surface (M, ℓ) satisfying Equation (5.4) at each edge is called *intrinsic Delaunay*.

Every polyhedral cone metric admits a Delaunay triangulation, *i.e.*, given a surface M with cone metric g , there always exists at least one triangulation of the cone points (vertices) where all edges satisfy Equation (5.4) [ILTC01]. In an extrinsic setting, one can imagine that a given polyhedron in \mathbb{R}^3 is triangulated by different collections of geodesic arcs terminating at vertices—not necessarily corresponding to the extrinsic edges (see especially Figure 15, right). Each such triangulation defines a discrete metric (given by the geodesic edge lengths), and one




FIGURE 15. A given polyhedron can be triangulated in many different ways; from an intrinsic point of view it does not matter whether there is a simplicial embedding into \mathbb{R}^n . For instance, the triangulations depicted at top and bottom both describe the same cone metric.




FIGURE 16. Left: an edge is Delaunay if the sum of opposite angles α, β is no greater than π ; “flipping” a non-Delaunay edge always yields a Delaunay configuration. Right: the angle sum condition holds with equality if and only if two triangles are cocircular.

such metric will always be Delaunay. A subtlety here is that this triangulation will not in general be simplicial—instead, one must consider *irregular* triangulations where, for example, two edges of the same triangle might be identified. (Such triangulations can be expressed as Δ -complexes, as discussed in Section 2).

As long as a Delaunay triangulation has no cocircular pairs of triangles, then it is also unique [BS07]. In other words, if two adjacent triangles ijk, jil share a common circumcircle when laid out in the plane, then either triangulation of the quadrilateral $ijkl$ will satisfy Equation (5.4) with equality (since opposite angles in a circular quadrilateral always sum to π). In general, replacing one triangulation of a quad $ijkl$ with the other is referred to as an *edge flip*. Given an initial triangulation, an intrinsic Delaunay triangulation can be obtained via a simple algorithm: flip non-Delaunay edges (in any order) until they all satisfy Equation (5.4); this algorithm will always terminate in a finite number of flips [BS07].

5.3. Discrete Uniformization. With several key pieces in place, one can now define a notion of discrete conformal equivalence for polyhedral surfaces with different combinatorics; this definition in turn leads naturally to a discrete uniformization theorem. There are two equivalent ways to state this definition, one which appeals only to Euclidean geometry, and another that makes use of a hyperbolic metric naturally associated with any Euclidean polyhedron (to be discussed in Part IV):

DEFINITION 5.6. Let (M, ℓ) and $(\tilde{M}, \tilde{\ell})$ be Delaunay triangulations of the same topological surface and with the same vertex set V . These triangulations are **discretely conformally equivalent** if either of two equivalent conditions hold:

- I. There is a sequence of Delaunay triangulations $(M, \ell) = (M_1, \ell_1), \dots, (M_n, \ell_n) = (\tilde{M}, \tilde{\ell})$ such that each pair of consecutive surfaces has either (i) identical combinatorics and equal length cross ratios (*i.e.*, they are related by a rescaling of edge lengths *à la* Equation (5.1)), or (ii) different combinatorics but identical Euclidean metric (*i.e.*, they are related by performing Euclidean edge flips on pairs of cocircular triangles).
- II. The ideal polyhedra associated with (M, ℓ) and $(\tilde{M}, \tilde{\ell})$ (as defined in Section 6.2) are related by a hyperbolic isometry.

This definition might seem somewhat limited, since at first glance it seems to apply “only” to Delaunay triangulations. Yet since every Euclidean polyhedron has

a Delaunay triangulation, Definition 5.6 is in fact much more natural than the one discussed in Section 5.1: it places emphasis on the geometry of the surface itself, rather than an arbitrary triangulation that merely sits on top of the surface.

Discrete uniformization then amounts to essentially the same statement as in the smooth setting: given any (discrete) metric, there is a (discretely) conformally equivalent one with constant curvature. More precisely, for a given discrete surface (M, ℓ) , there exists a conformally equivalent discrete metric with either constant cone angle Ω_i (Equation (2.1)), or more generally, that achieves some prescribed cone angle Ω_i^* (so long as it satisfies the discrete analogue of Gauss-Bonnet mentioned in Section 2). Several closely related versions of discrete uniformization are encapsulated by the following theorems.

THEOREM 5.7 (spherical uniformization). *For any closed finite discrete surface (M, ℓ) of genus zero, there exists a discretely conformally equivalent convex polyhedron $(\tilde{M}, \tilde{\ell})$ inscribed in the unit sphere $S^2 \subset \mathbb{R}^3$.*

THEOREM 5.8 (Euclidean uniformization). *For any closed finite genus-1 discrete surface (M, ℓ) with vertex set V , there exists a discretely conformally equivalent surface $(\tilde{M}, \tilde{\ell})$ with zero curvature at each vertex, i.e., $\Omega_i^* = 0$ for all $i \in V$.*

In the hyperbolic case (i.e., for genus $g \geq 2$), uniformization must be treated in a different way, since unlike the spherical case we cannot isometrically embed H^2 into \mathbb{R}^3 . Instead we endow our triangulation with a piecewise hyperbolic metric, this time constructed from ordinary rather than ideal hyperbolic triangles. In this setting, two discrete metrics $\ell, \tilde{\ell} : E \rightarrow \mathbb{R}_{>0}$ on the same triangulation $M = (V, E, F)$ are discretely conformally equivalent if they satisfy the relationship

$$\sinh(\tilde{\ell}_{ij}/2) = e^{(u_i+u_j)/2} \sinh(\ell_{ij}/2)$$

for each edge $ij \in E$ (for further discussion see [BPS15]). The definition of the cone angle Ω_i is unchanged: it is the difference between 2π and the angle sum Θ_i of the interior angles at vertex i . (This data also describes a Euclidean polyhedron inscribed in the hyperboloid model, where the Euclidean length of each edge is determined by the indefinite inner product $\langle x, x \rangle := x_1^2 + x_2^2 - x_3^2$.) The definition of conformal equivalence for variable triangulations is then directly analogous to Definition 5.6, yielding the following uniformization theorem:

THEOREM 5.9 (hyperbolic uniformization). *For any closed finite discrete surface (M, ℓ) of genus $g \geq 2$ with a piecewise hyperbolic metric, there exists a discretely conformally equivalent piecewise hyperbolic surface $(\tilde{M}, \tilde{\ell})$ with zero curvature at every vertex ($\Omega_i = 0$ for all $i \in V$).*

Theorem 5.8 is a special case of a more general uniformization theorem which guarantees existence of metrics with prescribed cone angles at each vertex:

THEOREM 5.10. *For any closed finite genus- g discrete surface (M, ℓ) with vertex set V , and any set of target cone angles $\Omega^* : V \rightarrow (-\infty, 2\pi)$ satisfying the Gauss-Bonnet condition*

$$\sum_{i \in V} \Omega_i^* = 2\pi(2 - 2g),$$

there exists a discretely conformally equivalent surface $(\tilde{M}, \tilde{\ell})$ with $\tilde{\Omega}_i = \Omega_i^$.*

From the Euclidean point of view, the basic idea behind proving discrete uniformization is to consider a time-continuous evolution of the discrete metric, akin to the same smooth Ricci/Yamabe flow discussed in Section 4.3. This flow was first studied by Luo in the context of a fixed triangulation [Luo04]. In this setting, the scale factors $u : V \rightarrow \mathbb{R}$ evolve according to the flow

$$(5.5) \quad \frac{d}{dt} u_i(t) = \Omega_i(t) - \Omega_i^*,$$

where $\Omega_i(t)$ is the angle defect corresponding to the current edge lengths $\tilde{\ell}(t) := e^{(u_i(t)+u_j(t))/2} \ell_{ij}$, and Ω_i^* is the (fixed) target curvature. Note that the curvatures $\Omega(t)$ depend on the angles θ , which in turn depend on the rescaled edge lengths $\ell_{ij}(t) = e^{(u_i(t)+u_j(t))/2} \ell_{ij}(0)$. Differentiating Equation (5.5) with respect to u reveals that the flow is the gradient flow on a convex energy $E(u)$, whose Hessian is given explicitly by the cotangent discretization of the Laplace-Beltrami operator Δ [Luo04]. An explicit form for this energy was first given by Springborn *et al.* [SSP08], and will be discussed in Section 6.4. Due to the convexity of the energy, the flow yields a uniformizing collection of scale factors $u^* : V \rightarrow \mathbb{R}$ in finite time—if such factors exist. However, for a fixed triangulation the flow may become singular, *i.e.*, it may reach a point where the edge lengths no longer describe a valid discrete metric (due to failure of the triangle inequality).

For this reason one must in general define discrete uniformization in terms of variable triangulations. Here one considers the same evolution of scale factors $u_i(t)$ as in Equation (5.5), but where the corresponding discrete metric $\tilde{\ell}(t)$ is always defined with respect to the intrinsic Delaunay triangulation of the current cone metric. At certain critical moments (namely, when there is a cocircular pair of triangles) this triangulation will not be unique, and effectively experiences an edge flip as the flow continues to evolve. One can continue to think of this flow as a gradient flow on a piecewise smooth energy whose definition is uniform over *Penner cells*, *i.e.*, regions of $\mathbb{R}^{|V|}$ over which the intrinsic Delaunay triangulation induced by the scale factors $u : V \rightarrow \mathbb{R}$ does not change. This energy remains C^1 (in fact C^2) everywhere—even at the boundary of such cells [Spr17]. One therefore has a well-defined gradient flow on a convex energy, and can then show that any discrete metric (with no special conditions on geometry or combinatorics) will flow to one of constant curvature in finite time.

The Euclidean case (Theorem 5.8 and Theorem 5.10), was first proven by Gu *et al.* [GLSW18]; the hyperbolic case (Theorem 5.9) is implied by a proof by Fillastre [Fil08] about hyperbolic polyhedra, and was shown more directly by Gu *et al.* [GGL⁺18]. The spherical case (Theorem 5.7) was first shown by Springborn [Spr17], via a proof that also covers the Euclidean and hyperbolic cases. None of these proofs consider domains with boundary, nor conditions on monodromy around noncontractible cycles [CZ17]. In order for a uniformization to be algorithmically computable, it may also be important (depending on the choice of algorithm) that the number of edge flips encountered during the flow is finite, *i.e.*, that the flow passes through only finitely many Penner cells; this fact was established by Wu [Wu14].

Convergence under mesh refinement of discrete conformal maps to smooth ones has been studied both numerically and analytically [SWGL15, Büc16, Büc18a] (see also [SC17, Figure 16]), though many questions still remain.

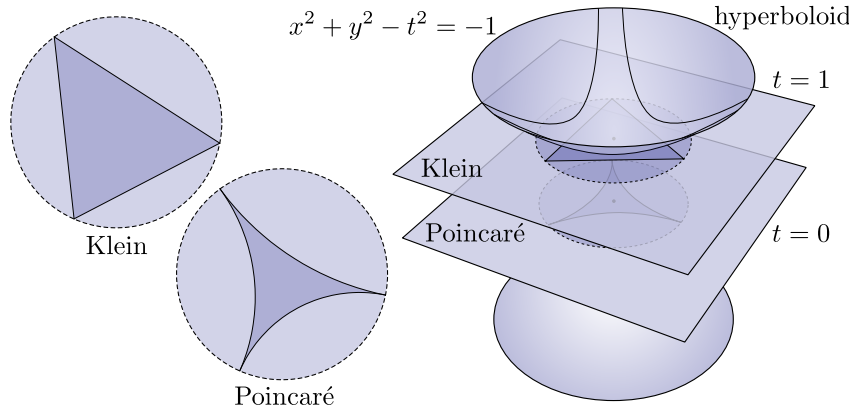



FIGURE 17. Left: an ideal triangle in two models of the hyperbolic plane H^2 , which preserve either straight lines (Klein) or angles (Poincaré). Right: these models are both related to the hyperboloid model via projections onto the planes $t = 0$ and $t = 1$ (*resp.*).

6. Part IV: Hyperbolic Polyhedra

In this final part we consider the hyperbolic viewpoint on discrete conformal geometry, which helps to simplify and unify the story. Along the way we will also encounter some fascinating connections between geometry and combinatorics, which instigated the development of techniques now used for discrete conformal geometry.

Why is hyperbolic geometry an effective framework for problems in discrete conformal geometry? For one thing, ideal hyperbolic polyhedra are often “easier” to construct than Euclidean ones, since they typically exhibit greater rigidity. For instance, the dihedral angles of a convex polyhedron do not determine its extrinsic geometry in the Euclidean case (consider, for instance, a family of truncated polyhedra), whereas convex hyperbolic polyhedra are essentially determined by their dihedral angles [HR93]. Moreover, even though Alexandrov’s theorem guarantees that a convex Euclidean polyhedron is uniquely determined by its metric (up to rigid motions), actually constructing an embedding is quite challenging [IB108, KPD09]. In contrast, an embedding of an ideal convex hyperbolic polyhedron with prescribed metric can be obtained as the minimizer of a convex energy [Spr17]. These two realization problems—ideal polyhedra with prescribed dihedral angles or prescribed metric—turn out to be closely connected (and in some cases, identical) to the problem of finding circle patterns with prescribed intersection angles, or discretely conformally equivalent metrics with prescribed cone angles, *resp.*



In order to describe these connections, we first give some background on hyperbolic geometry, followed by a discussion of variational principles for hyperbolic polyhedra and their connection to discrete conformal maps.

6.1. Hyperbolic Geometry. The hyperbolic plane H^2 is topologically like the ordinary Euclidean plane, but has “saddle-like” geometry everywhere; more

formally, it is a complete and simply-connected surface of constant negative curvature $K = -1$. Unlike the unit 2-sphere ($K = +1$), the hyperbolic plane cannot be isometrically embedded in \mathbb{R}^3 and must instead be studied in terms of various *models*. Two models that naturally arise in discrete conformal geometry are the *Poincaré disk model* and the *Klein disk model*, which both model H^2 on the open unit disk $D^2 \subset \mathbb{R}^2$, as depicted in Figure 17, *left*. The most salient geometric facts about these two models is that the Poincaré model is conformal and hence faithfully represents angles but not straight lines (geodesics), whereas the Klein model faithfully represents geodesics as straight lines but does not preserve angles. In the Poincaré model, geodesics instead become circular arcs orthogonal to the unit circle ∂D^2 . These two models can be connected via a third *hyperboloid model*, which puts H^2 into correspondence with one sheet of the hyperboloid $x^2 + y^2 - t^2 = -1$; in this model, geodesics are represented as intersections of the hyperboloid with planes through the origin. If we express points on the hyperboloid in homogeneous coordinates (x, y, t) , then the map from the hyperboloid to the Klein model is given by the usual homogeneous projection $(x, y, t) \mapsto (x/t, y/t)$. Similarly, to get from the hyperboloid model to the Poincaré model we can simply shift the hyperboloid a unit distance along the t -axis before applying the same projection. These relationships are depicted in Figure 17, *right*.

Unlike the Euclidean plane, where a pair of parallel lines remains at a constant distance, one can construct pairs of hyperbolic geodesics that never intersect yet become arbitrarily close as they approach infinity (such pairs are sometimes called *limiting parallels*). An *ideal hyperbolic triangle* is a figure bounded by three such pairs; its *vertices* are the three corresponding limit points on the sphere at infinity. As a result, all interior angles are zero and all edges have infinite length. In fact, all ideal triangles are *congruent*, *i.e.*, they are identical up to isometries of H^2 , which are represented in the Poincaré model and the Klein model as Möbius transformations and projective transformations (*resp.*) that map the disk to itself.

6.2. Ideal Hyperbolic Polyhedra. Any Euclidean triangulation can naturally be identified with one made of ideal hyperbolic triangles: take each triangle and draw it in the Euclidean plane. The circumcircle of this triangle (*i.e.*, the unique circle passing through its three vertices) can be viewed as a copy of the hyperbolic plane in the Klein model. The three straight edges then correspond to three hyperbolic geodesics, and the triangle itself becomes an ideal hyperbolic triangle. By gluing these ideal triangles together along shared edges, we obtain an *ideal hyperbolic polyhedron*, *i.e.*, a surface of constant curvature $K = 1$ away from a collection of *cusps* at vertices, which extend to infinity (examples are shown in Figure 18 and Figure 22).

Since all ideal triangles are congruent, it would be easy to think that the geometry of such a polyhedron is determined purely by the combinatorics of M . However, additional structure is encoded in the way pairs of ideal triangles are identified: given two ideal triangles ijk, jil sharing an edge γ , we can apply a hyperbolic isometry to one of them (say, jil) that fixes γ , *i.e.*, we can “slide” one ideal triangle along the other. The intrinsic geometry of the hyperbolic polyhedron is therefore fully determined by (i) the combinatorics of the triangulation, and (ii) data that specifies how edges are identified.





FIGURE 18. A Euclidean polyhedron and its realization as an ideal hyperbolic polyhedron.

To explicitly encode identifications, Penner [Pen87] defines *shear coordinates* $\sigma_{ij} \in \mathbb{R}$ per edge $ij \in E$ which measure the distance between the altitudes of the two triangles with base ij , as depicted in the inset. (Note that these distances are measured with respect to the hyperbolic metric, and have nothing to do with a particular model or choice of coordinates on H^2 .) Alternatively, one can “decorate” the vertices of the ideal polyhedron with arbitrarily-chosen *horocycles*, which are copies of Euclidean \mathbb{R}^n sitting inside H^n . Just as a Euclidean line can be viewed as a circle of infinite radius, a horocycle can



be viewed as the limit of a family of mutually tangent hyperbolic circles as the radius goes to infinity—in the Poincaré model, any such curve is represented by a circle tangent to the boundary. For a triangle ijk , the horocycle associated with a given vertex i will meet the two hyperbolic geodesics corresponding to edges ij and ik orthogonally. The *Penner coordinate* $\lambda_{ij} \in \mathbb{R}$ of each edge ij is then the (hyperbolic) distance between the horocycles at i and j . More precisely, λ_{ij} gives the *signed* distance between horocycles: positive if they are disjoint, negative if they intersect. Since the choice of horocycles is arbitrary, the Penner coordinate of a *single* edge does not, in isolation, provide any geometric information about the hyperbolic structure of the polyhedron. However, the Penner coordinates share an important relationship with the shear coordinates, namely

$$(6.1) \quad 2\sigma_{ij} = \lambda_{jk} - \lambda_{ik} + \lambda_{il} - \lambda_{jl}.$$

Hence, two sets of Penner coordinates determine the same ideal hyperbolic polyhedron if and only if they induce the same shear coordinates.

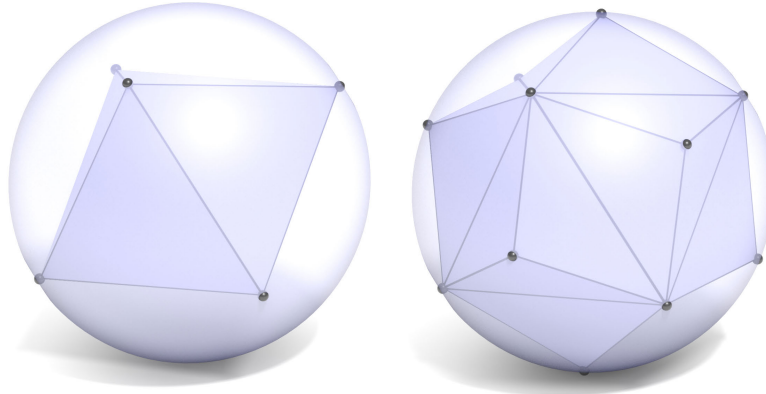


FIGURE 19. Unlike the octahedron (left), the stellated octahedron (right) cannot be realized a convex spherical polyhedron, no matter where the vertices are placed.

The similarity between Equation (6.1) and the definition of the length cross ratio (Equation (5.3)) is not superficial: when an ideal hyperbolic polyhedron is constructed from a Euclidean polyhedron (as described above), the shear coordinates and length cross ratios will be related by

$$\sigma_{ij} = \log(\mathbf{c}_{ij}),$$

independent of any choice of horocycles. It is therefore natural to pick horocycles such that

$$\lambda_{ij} = 2 \log \ell_{ij},$$

i.e., such that the Penner coordinates encode the edge lengths. From the conformal point of view, the shear coordinates encode the discrete conformal equivalence class, and the Penner coordinates specify a particular discrete metric within this class.

Importantly, since both the shear coordinates and Penner coordinates are hyperbolic distances, they will be preserved by hyperbolic isometries. Hence, *isometry classes of ideal hyperbolic polyhedra correspond to conformal equivalence classes of discrete metrics*. This basic viewpoint was introduced by Bobenko, Pinkall, and Springborn; see in particular [BPS15, Section 5] for further discussion.

6.3. Variational Principles. A *variational principle* expresses the solution to a problem as a minimizer of an energy functional (possibly subject to constraints). A major appeal of discrete conformal geometry is that the variational principles involved are often *convex*, providing a clear picture of existence and uniqueness, and allowing one to leverage principled methods from convex optimization [BV04] to develop practical algorithms with clear guarantees. Much of the understanding of these principles originated in the hyperbolic setting, where they can be given a concrete geometric interpretation. Here, the basic question is not how to uniformize a Euclidean polyhedron, but rather how to construct embeddings of hyperbolic polyhedra with prescribed geometric data—namely, dihedral angles or the intrinsic metric.

A Prelude: Steiner’s Problem. Why is anything known about hyperbolic polyhedra in the first place? One answer is that triangulations of hyperbolic space are a key tool for understanding the geometry and topology of 3-manifolds [Thu98]. A more basic question was posed by the 19th century Swiss geometer Jakob Steiner, which provides a first glimpse at some of the deep connections between geometry and combinatorics:

“Wenn irgend ein convexes Polyöder gegeben ist, läßt sich dann immer (oder in welchen Fällen nur) irgend ein anderes, welches mit ihm in Hinsicht der Art und der Zusammensetzung der Grenzflächen übereinstimmt (oder von gleicher Gattung ist), in oder um eine Kugelfläche, oder in oder um irgend eine andere Fläche zweiten Grades beschreiben (d. h. daß seine Ecken alle in dieser Fläche liegen, oder seine Grenzflächen alle diese Fläche berühren)?”

In English: “Given a convex polyhedron, when can you find a combinatorially equivalent convex polyhedron that is inscribed in the sphere (or another quadratic surface)?” More abstractly, we might ask: which combinatorial tessellations of the 2-sphere can be realized as convex polyhedra with vertices on the sphere? Steiner’s question is very much in the spirit of discrete differential geometry: he is looking for a finite analogue of the sphere that exactly preserves a key property of its smooth counterpart (convexity).

Surprisingly enough, not all tessellations of the sphere can be realized as convex spherical polyhedra: for instance the octahedron sits in the sphere, but for the *stellated octahedron* (where we split each face into three) there turns out to be no way to arrange the vertices on the sphere so that it becomes convex [Riv93, Corollary 10]. In other words, only certain “nice” combinatorial tessellations of the topological sphere can be interpreted as geometric spheres. This same question can also be stated in terms of Delaunay triangulations: every planar Delaunay triangulation is the image of some sphere-inscribed convex polyhedron under stereographic projection. Hence, Steiner’s question amounts to asking which combinatorial triangulations can be realized as planar Delaunay triangulations.

More than 150 years later, a general solution to Steiner’s problem was found by framing it as a problem about hyperbolic rather than Euclidean polyhedra [Riv96]. The bridge is the one we have already seen (in Section 6.2): Euclidean polyhedra determine a hyperbolic structure by interpreting each Euclidean triangle as an ideal triangle in the Klein model; constructing a convex sphere-inscribed Euclidean polyhedron with prescribed combinatorics then becomes equivalent to finding a convex ideal hyperbolic polyhedron with the same combinatorics. As we will see, this problem turns out to be closely linked to the variational principles for both the circle patterns studied in Part II, and the conformally equivalent discrete metrics studied in Part III.

Lobachevsky’s Function. Perhaps the most concrete link between hyperbolic and discrete conformal geometry is through the *Lobachevsky function*

$$\mathbb{L}(\theta) := - \int_0^\theta \log |2 \sin u| \, du,$$

Thurston provides an account of some of its analytical and geometric properties [Thu79, Chapter 7]. Though originally developed for calculating hyperbolic




FIGURE 20. Any Euclidean triangle ijk with interior angles θ can be identified with an ideal hyperbolic tetrahedron $ijkl$ with identical dihedral angles. The logarithmic edge lengths $\lambda_{ij} = 2 \log \ell_{ij}$ determine a choice of horospheres at vertices.

volumes, this function also arises naturally in variational principles for discrete conformal maps.

In particular, consider an *ideal hyperbolic tetrahedron*, *i.e.*, an ideal hyperbolic polyhedron with the usual combinatorics of a tetrahedron. One can show that the volume of such a tetrahedron is

$$V(\alpha) := \frac{1}{2} \sum_{ij} \mathbb{I}(\alpha_{ij}),$$

where the sum is taken over all six edges, and α_{ij} is the dihedral angle associated with edge ij (for a proof, see [Thu79, Theorem 7.2.1]). Since opposite dihedral angles in an ideal tetrahedron are equal, only three of the dihedral angles are distinct. Moreover, these angles, which we will call α_{ij} , α_{jk} , and α_{ki} , can be naturally identified with the interior angles θ_k^{ij} , θ_i^{jk} , θ_j^{ki} of a Euclidean triangle (*resp.*), as depicted in Figure 20. Importantly, then, the function $-V$ is convex when restricted to the set of valid Euclidean angles, *i.e.*, whenever $\alpha_{ij} + \alpha_{jk} + \alpha_{ki} = \pi$ and all the angles are positive [Riv94, Theorem 2.1]. The first order change in volume \dot{V} with respect to a first order change $\dot{\alpha}$ in dihedral angles can be expressed via the following *Schläfli formula* for ideal tetrahedra:

$$(6.2) \quad \dot{V} = -\frac{1}{2} \sum_{ij} \lambda_{ij} \dot{\alpha}_{ij},$$

where λ_{ij} are the Penner coordinates of edge ij [Mil94]; Rivin provides a proof [Riv94, Theorem 14.5].




FIGURE 21. In the Poincaré halfspace model for H^3 , two Euclidean circles intersecting at an angle Φ_{ij} can be interpreted as two copies of H^2 intersecting at the same dihedral angle. Three such intersections trace out an ideal hyperbolic triangle.

Such observations provide a starting point for convex variational principles for hyperbolic polyhedra [HR93, Riv93, Riv96, Fil08], which in turn provide variational principles for discrete conformal geometry. Consider, for instance, a Euclidean circle pattern associated with the faces of a triangulation \mathbf{M} , where circles associated with neighboring triangles ijk, jil intersect at angles Φ_{ij} [BS04a]. If we view each circle as the boundary of a hemisphere sitting on top of the plane, then the intersection angles coincide with the dihedral angles between neighboring spheres. In the Poincaré half space model for H^3 , each such hemisphere corresponds to a copy of H^2 . The intersections between spheres are then hyperbolic geodesics bounding ideal hyperbolic triangles. Hence, circle patterns with prescribed intersection angles are equivalent to convex ideal polyhedra with prescribed dihedral angles, which are in turn equivalent to convex Euclidean polyhedra inscribed in the sphere (via the Klein model, as discussed in Section 6.2), which are in turn equivalent to planar Delaunay triangulations, via stereographic projection. The connections between these different points of view are summarized in Figure 22.

Alternatively, suppose we draw an ideal tetrahedron in the Poincaré half space model, as depicted in Figure 20. Then the three distinct dihedral angles α_{ij} correspond to the three interior angles of a Euclidean triangle (at infinity), and we can always pick horospheres at vertices such that the Penner coordinates λ_{ij} are determined by the logarithmic edge lengths, *i.e.*, such that $\ell_{ij} = e^{\lambda_{ij}/2}$. By Equation (6.2), the first order variation of the function

$$(6.3) \quad \varphi(\alpha, x) := V(\alpha) + \sum_{ij} \lambda_{ij} \alpha_{ij}$$

is then given by

$$\dot{\varphi} = \dot{V} + \sum_{ij} \dot{\lambda}_{ij} \alpha_{ij} + \sum_{ij} \lambda_{ij} \dot{\alpha}_{ij} = \sum_{ij} \dot{\lambda}_{ij} \alpha_{ij}.$$

Hence,

$$\frac{\partial \varphi}{\partial \lambda_{ij}} = \alpha_{ij},$$

allowing us to obtain the interior angles of a triangle as the derivatives of a convex functional φ (see also [BPS15, Section 4.2]). This functional therefore provides a starting point for variational principles for problems involving an evolution of a discrete metric $\ell = e^{\lambda/2}$. In particular, it leads to a variational principle for discrete

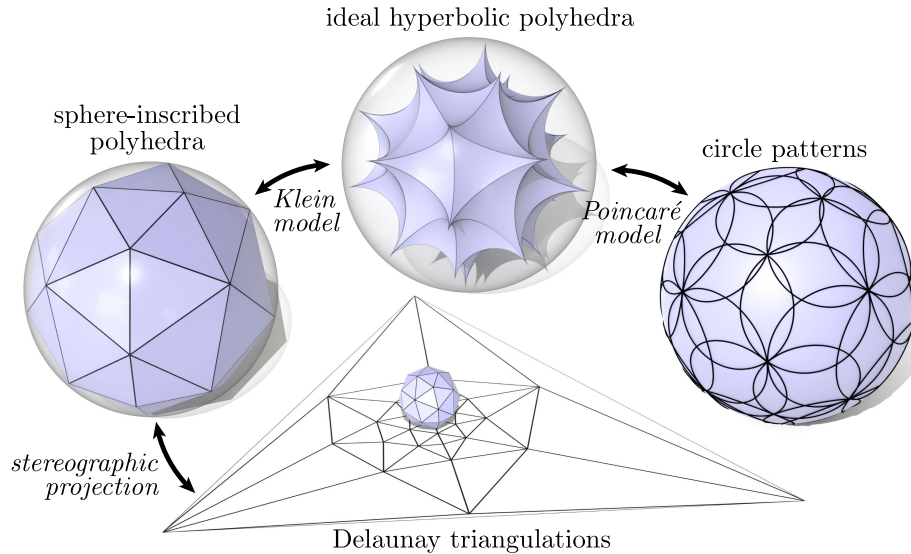


FIGURE 22. A triangulation with fixed combinatorics can be realized as a Delaunay triangulation in the plane if and only if it can also be realized as a convex Euclidean polyhedron inscribed in the sphere, a convex ideal hyperbolic triangulation, and a circle pattern on the sphere (with circles associated with triangles).

uniformization involving the cone angles Ω_i (Section 6.4), which can be expressed in terms of sums of interior angles θ , or equivalently, dihedral angles α .

Chronologically, the presentation in this section is backwards: variational principles in discrete conformal geometry were historically developed by first establishing some kind of “flow” involving only derivatives (such as the combinatorial Ricci flow discussed in Section 4.3); only later were these derivatives integrated to obtain explicit expressions for the underlying *variational principle*, *i.e.*, the energy driving the flow. While derivatives alone are often enough to establish results about existence, uniqueness, convergence, *etc.*, variational principles are quite valuable for understanding the global picture (including connections between Euclidean and hyperbolic problems) as well as for practical algorithms (enabling, for instance, the use of principled optimization strategies [KSS06, SSP08]).

Circle Patterns. Early approaches to constructing circle packings can be loosely interpreted as iterative algorithms (*e.g.*, Jacobi or Gauss-Seidel) for minimizing an energy [Thu79, CS03]. Colin de Verdière [Col91] gave the first real variational principle for Euclidean and hyperbolic circle patterns with circles at vertices, though only expressions for the derivatives were given; Brägger [Bra92] makes a very brief remark that these derivatives appear related to hyperbolic volume (“*Untersucht man diese Funktionale, stösst man rasch auf die Lobachevsky-Funktion und damit auf das Volumen hyperbolischer 3-Simplices.*”). Rivin [Riv94] considers ideal polyhedra with prescribed dihedral angles, effectively providing a variational principle for Euclidean circle patterns. Bobenko & Springborn [BS04a] develop variational principles that connect and generalize the three previously mentioned, and leads to efficient numerical implementation [KSS06] due to the absence of

difficult constraints. Chow and Luo [CL03] subsequently introduce a combinatorial analogue of Ricci flow (Section 4.3), and show that it arises from a convex variational principle by considering the Hessian [CL03, Section 3.5]. More recently, Ge also considers a combinatorial analogue of Calabi flow [Ge18] suitable for practical computation [ZLH⁺18]. A variety of closely related variational principles and generalizations thereof have also been studied [Gli11, ZGZ⁺14, GX15], including a large number of applications [JKG07].

Conformally Equivalent Discrete Metrics. The first connection between variational principles and conformally equivalent discrete metrics (*à la* Definition 5.1) came from Roček and Williams’ efforts to study quantum gravity through the simplicial calculus of Regge [Reg61]. The definition of discrete conformal equivalence here is identical to the one discussed in Section 5.1 [RW84, Section VII], though the discrete variational principles are mainly connected to Einstein metrics on four dimensional spacetime rather than metrics of constant scalar curvature on 2-manifolds. Luo [Luo04] later developed the same notion of conformally equivalent discrete metrics, and a corresponding discrete Yamabe flow driving the metric toward constant scalar curvature; no explicit energy was given, and existence is guaranteed only if the flow does not develop singularities (*i.e.*, if the discrete metric does not degenerate). Springborn, Schröder, and Pinkall gave an explicit energy in terms of the Lobachevsky function [SSP08], and Bobenko, Pinkall, and Springborn [BPS15] made the connection to hyperbolic geometry and ideal polyhedra. Through this correspondence, Fillastre effectively proved discrete uniformization in the hyperbolic case [Fil08]; Gu *et al.* establish the same result from the perspective of discrete conformal equivalence [GGL⁺18, Corollary 4 and Theorem 5]. Gu *et al.* also establish the existence of discrete uniformization in the Euclidean setting [GLSW18]; Springborn completes the proof of discrete uniformization via a variational approach that subsumes the Euclidean, hyperbolic, and spherical case [Spr17].

Much more has been said about variational principles both for ideal hyperbolic polyhedra, and for polyhedra in general—see for instance [BS04a, DGL08].

6.4. Variational Principle for Discrete Uniformization. The hyperbolic picture provides an attractive perspective on the discrete uniformization theorems discussed in Section 5.3. Recall that two Euclidean polyhedra are discretely conformally equivalent if and only if their Delaunay triangulations induce the same piecewise hyperbolic metric (Definition 5.6). Intuitively, the reason for using a canonical triangulation like the intrinsic Delaunay triangulation is that different triangulations of the same Euclidean polyhedron will in general induce a *different* hyperbolic metric. The Delaunay triangulation is particularly special because Euclidean and hyperbolic Delaunay flips coincide exactly when the triangulation is not unique, *i.e.*, when two adjacent triangles are cocircular. Consider for instance two triangles $ijk, jil \in \mathbb{F}$, realized as a pair of ideal hyperbolic triangles in the Klein model. In this setting, the new edge length ℓ_{kl} resulting from a flip is determined by *Ptolemy’s relation*:

$$(6.4) \quad \ell_{ij}\ell_{kl} = \ell_{ki}\ell_{lj} + \ell_{il}\ell_{jk}.$$

Such a flip is therefore called a *Ptolemy flip*, and preserves the hyperbolic structure of M , *i.e.*, the new and old hyperbolic metric are the same. However, the new and old *Euclidean* metric will be different except in the special case where ijk and jil have the same circumcircle, since in this case the new edge length ℓ_{kl} resulting

from a Euclidean flip or a Ptolemy flip will be identical. This fact is somewhat counterintuitive: it says that even for a completely flat domain, a Euclidean edge flip will (in general) change the discrete conformal structure—even though it does not change the Euclidean geometry. Much as a Euclidean edge flip preserves the metric only if the edge is flat (*i.e.*, has zero dihedral angle), a hyperbolic edge flip will preserve the metric only if the hyperbolic edge has zero dihedral angle, or equivalently, if the four vertices are cocyclic. The Delaunay condition is therefore essential for the equivalence of definitions (I) and (II): since all meshes in the sequence are Delaunay, any Euclidean edge flip will also correspond to a Ptolemy flip (and hence preserve the hyperbolic metric). Alternatively one could also say: two triangulations are discretely conformally equivalent if they are related by a sequence of length scalings (*à la* Equation (5.1)) and *Ptolemy* flips, in which case the intermediate triangulations are no longer required to be Delaunay.

The problem of finding a discretely conformally equivalent flat metric (*i.e.*, all cone angles equal to zero) is then the same as finding an ideal polyhedron in H^3 with prescribed hyperbolic metric. Both problems can be solved by considering the same variational principle, which is discussed at length in a sequence of papers [SSP08, BPS15, Spr17]. In particular, let ℓ be an initial discrete metric on a triangulation M , and for any conformally equivalent metric $\tilde{\ell}$ with scale factors u at vertices, define the following energy with respect to the intrinsic Delaunay triangulation of $\tilde{\ell}$:

$$E(u) := \sum_{ijk \in F} \varphi(\tilde{\lambda}_{ij}, \tilde{\lambda}_{jk}, \tilde{\lambda}_{ki}) - \frac{\pi}{2}(u_i + u_j + u_k) + \frac{1}{2} \sum_{i \in V} (2\pi - \Omega_i^*) u_i.$$

Here, φ is the potential given in Equation (6.3), $\tilde{\lambda}_{ij} = 2 \log(\tilde{\ell}_{ij})$ are the logarithmic edge lengths, and Ω^* are the target cone angles [SSP08]. One can show that this energy is convex, and in fact C^2 everywhere [Spr17]. To uniformize a given discrete metric over a Euclidean domain (possibly with cone singularities), one can therefore apply standard first- or second-order descent to E with respect to the log conformal factors u ; the energy simply needs to be expressed relative to the appropriate intrinsic Delaunay triangulation, which can always be obtained by applying a sequence of Ptolemy flips. This principle (and variations of it) provide one approach to proving the discrete uniformization theorems discussed in Section 5.3 [Spr17].

7. Summary

The table below summarizes several different approaches to the discretization of conformal maps. To date, the approach based on conformally equivalent discrete metrics appears to be the only one that clearly exhibits both the right amount of flexibility, *and* furnishes a complete discrete uniformization theorem. However, the perspective of circle preservation comes very close—and ultimately it should not be too surprising to also find a complete picture of discrete uniformization in terms of circles. In fact, although these two perspectives at first appear somewhat disjoint, connections continue to be found at various levels. For instance, Glickenstein [Gli11] and Bücking [Büc18b] discuss connections between circle patterns and length cross ratios; Lam and Pinkall [LP16] show that linearized versions of both theories can be connected to discrete harmonic functions (in the sense of Equation (3.4)). More elementary starting points, natural as they may seem, lead to definitions that turn out to be too rigid or too flexible. However, as noted in Section 1 and discussed

throughout, many of these discretizations nonetheless provide interesting connections to the smooth theory, and play an important role in practical algorithms.

Beyond the basic question of flattening/uniformization, many questions about how to discretize conformal geometry remain. For instance, questions about *extrinsic* conformal maps, *i.e.*, transformations of conformal surface immersions in \mathbb{R}^n , are relevant for practical tasks in digital geometry processing, where one wishes to directly manipulate geometry sitting in space, rather than mapping it to a canonical domain. Some preliminary work has been done in this direction [Cra13], including a recent theory compatible with the notion of conformally equivalent discrete metrics [LP18]. There has also been some recent work on applying similar tools to the problem of discrete *geometrization* [AMY18], as well as continued interest in problems in numerical relativity [Gen02] via Regge calculus.

Approach	Data	Outcome	Comments
ANGLES (Setion 3.1)	interior angles θ_i^{jk}	too rigid	similarity of triangles forces single global scale factor
DIRICHLET (Setion 3.3)	vertex coordinates f_i	too rigid	(same as ANGLES)
HODGE (Setion 3.4)	length ratio w_{ij}	too rigid	uniquely determines a discrete metric
CONJUGATE (Setion 3.5)	vertex coordinates f_i	just right	<i>only under refinement / no finite notion of conformal equivalence</i>
CIRCLES (Setion 4)	graph $G = (V, E)$	too flexible	only combinatorics are considered; no way to distinguish different metrics
	intersection angles α_{ij}	just right	<i>existence not guaranteed for all triangulations</i>
METRIC (Setion 5)	edge lengths ℓ_{ij}	just right	existence is guaranteed

Acknowledgements. Thanks to Feng Luo and Boris Springborn for answering questions about discrete uniformization, and to the anonymous referees for providing additional useful comments and references.

References

- [AMY18] Paul M. Alsing, Warner A. Miller, and Shing-Tung Yau, *A realization of Thurston’s geometrization: discrete Ricci flow with surgery*, Ann. Math. Sci. Appl. **3** (2018), no. 1, 31–45. MR 3781260
- [BCGB08] Mirela Ben-Chen, Craig Gotsman, and Guy Bunin, *Conformal Flattening by Curvature Prescription and Metric Scaling*, CG Forum (2008).
- [BCK18] Alex Baden, Keenan Crane, and Misha Kazhdan, *Möbius Registration*, Computer Graphics Forum (SGP) **37** (2018), no. 5.
- [BDS17] Alexander I. Bobenko, Nikolay Dimitrov, and Stefan Sechelmann, *Discrete uniformization of polyhedral surfaces with non-positive curvature and branched covers over the sphere via hyper-ideal circle patterns*, Discrete Comput. Geom. **57** (2017), no. 2, 431–469. MR 3602861
- [BG17] Alexander Bobenko and Felix Günther, *Discrete Riemann Surfaces Based on Quadrilateral Cellular Decompositions*, Adv. Math. **311** (2017), 885–932.
- [BHS06] Alexander I. Bobenko, Tim Hoffmann, and Boris A. Springborn, *Minimal surfaces from circle patterns: geometry from combinatorics*, Ann. of Math. (2) **164** (2006), no. 1, 231–264. MR 2233848

- [BPS15] Alexander I. Bobenko, Ulrich Pinkall, and Boris A. Springborn, *Discrete conformal maps and ideal hyperbolic polyhedra*, *Geom. Topol.* **19** (2015), no. 4, 2155–2215. MR 3375525
- [Bra92] W Bragger, *Kreispackungen und Triangulierungen*, *Enseign. Math.* **38** (1992), 201–217.
- [BS04a] Alexander I. Bobenko and Boris A. Springborn, *Variational principles for circle patterns and Koebe’s theorem*, *Trans. Amer. Math. Soc.* **356** (2004), no. 2, 659–689. MR 2022715
- [BS04b] Philip Bowers and Kenneth Stephenson, *Uniformizing dessins and belyi maps via circle packing*, vol. 170, American Mathematical Soc., 2004.
- [BS07] Alexander I. Bobenko and Boris A. Springborn, *A discrete laplace–beltrami operator for simplicial surfaces*, *Discrete & Computational Geometry* **38** (2007), no. 4, 740–756.
- [Büc16] Ulrike Bücking, *Approximation of conformal mappings using conformally equivalent triangular lattices*, *Advances in Discrete Differential Geometry* (Alexander I. Bobenko, ed.), Springer Berlin Heidelberg, Berlin, Heidelberg, 2016, pp. 133–149.
- [Büc18a] ———, *C^∞ -Convergence of Conformal Mappings for Conformally Equivalent Triangular Lattices*, *Results in Mathematics* **73** (2018), no. 2, 84.
- [Büc18b] ———, *Conformally symmetric triangular lattices and discrete ϑ -conformal maps*, arXiv e-prints (2018), arXiv:1808.08064.
- [BV04] Stephen Boyd and Lieven Vandenbergh, *Convex optimization*, Cambridge University Press, New York, NY, USA, 2004.
- [CdGDS13] Keenan Crane, Fernando de Goes, Mathieu Desbrun, and Peter Schröder, *Digital geometry processing with discrete exterior calculus*, ACM SIGGRAPH 2013 Courses, SIGGRAPH ’13, ACM, 2013.
- [Cer96] Davide P Cervone, *Tight immersions of simplicial surfaces in three space*, *Topology* **35** (1996), no. 4, 863 – 873.
- [CL03] Bennett Chow and Feng Luo, *Combinatorial ricci flows on surfaces*, *J. Differential Geom.* **63** (2003), no. 1, 97–129.
- [Col91] Yves Colin de Verdière, *Un principe variationnel pour les empilements de cercles*, *Invent. Math.* **104** (1991), no. 3. MR 1106755 (92h:57020)
- [Cra13] Keenan Crane, *Conformal geometry processing*, Ph.D. thesis, Caltech, June 2013.
- [CS03] Charles R. Collins and Kenneth Stephenson, *A circle packing algorithm*, *Computational Geometry* **25** (2003), no. 3, 233 – 256.
- [CSD02] David Cohen-Steiner and Mathieu Desbrun, *Hindsight: LSCM and DNCP are One and the Same*, 2002.
- [CW17] Keenan Crane and Max Wardetzky, *A glimpse into discrete differential geometry*, *Notices of the American Mathematical Society* **64** (2017), no. 10, 1153–1159.
- [CZ17] M. Campen and D. Zorin, *On Discrete Conformal Seamless Similarity Maps*, ArXiv e-prints (2017).
- [DGL08] J. Dai, X.D. Gu, and F. Luo, *Variational principles for discrete surfaces*, *Advanced lectures in mathematics*, International Press, 2008.
- [DHLM05] M. Desbrun, A. N. Hirani, M. Leok, and J. E. Marsden, *Discrete Exterior Calculus*, ArXiv Mathematics e-prints (2005).
- [DMA02] Mathieu Desbrun, Mark Meyer, and Pierre Alliez, *Intrinsic parameterizations of surface meshes*, *Computer Graphics Forum* **21** (2002), no. 3, 209–218.
- [Fil08] François Fillastre, *Polyhedral hyperbolic metrics on surfaces*, *Geometriae Dedicata* **134** (2008), no. 1, 177–196.
- [Ge18] Huabin Ge, *Combinatorial Calabi flows on surfaces*, *Trans. Amer. Math. Soc.* **370** (2018), no. 2, 1377–1391. MR 3729504
- [Gen02] Adrian P. Gentle, *Regge calculus: a unique tool for numerical relativity*, *Gen. Relativity Gravitation* **34** (2002), no. 10, 1701–1718. MR 1938626
- [GGL⁺18] Xianfeng Gu, Ren Guo, Feng Luo, Jian Sun, and Tianqi Wu, *A discrete uniformization theorem for polyhedral surfaces II*, *J. Differential Geom.* **109** (2018), no. 3, 431–466. MR 3825607
- [Gli11] David Glickenstein, *Discrete conformal variations and scalar curvature on piecewise flat two- and three-dimensional manifolds*, *J. Differential Geom.* **87** (2011), no. 2, 201–237. MR 2788656

- [GLSW18] Xianfeng David Gu, Feng Luo, Jian Sun, and Tianqi Wu, *A discrete uniformization theorem for polyhedral surfaces*, J. Differential Geom. **109** (2018), no. 2, 223–256. MR 3807319
- [GMMD14] Fernando de Goes, Pooran Memari, Patrick Mullen, and Mathieu Desbrun, *Weighted triangulations for geometry processing*, ACM Trans. Graph. **33** (2014), no. 3, 28:1–28:13.
- [Guo11] Ren Guo, *Local rigidity of inversive distance circle packing*, Trans. Amer. Math. Soc. **363** (2011), no. 9, 4757–4776. MR 2806690
- [GX15] H. Ge and X. Xu, *A Discrete Ricci Flow on Surfaces in Hyperbolic Background Geometry*, ArXiv e-prints (2015).
- [GY08] X.D. Gu and S.T. Yau, *Computational conformal geometry*, Advanced lectures in mathematics, International Press, 2008.
- [Ham82] Richard S. Hamilton, *Three-manifolds with positive ricci curvature*, J. Differential Geom. **17** (1982), no. 2, 255–306.
- [Hat02] A. Hatcher, *Algebraic topology*, Algebraic Topology, Cambridge University Press, 2002.
- [HK15] J. Hass and P. Koehl, *A Metric for genus-zero surfaces*, arXiv e-prints (2015).
- [HR93] Craig D. Hodgson and Igor Rivin, *A characterization of compact convex polyhedra in hyperbolic 3-space*, Inventiones mathematicae **111** (1993), no. 1, 77–111.
- [HS15] Joel Hass and Peter Scott, *Simplicial energy and simplicial harmonic maps*, Asian J. Math. **19** (2015), no. 4, 593–636. MR 3423736
- [Hut91] John E. Hutchinson, *Computing conformal maps and minimal surfaces*, Theoretical and Numerical Aspects of Geometric Variational Problems (Canberra AUS), Centre for Mathematics and its Applications, Mathematical Sciences Institute, The Australian National University, 1991, pp. 140–161.
- [IBI08] Alexander I. Bobenko and Ivan Izmistiev, *Alexandrov’s theorem, weighted delaunay triangulations, and mixed volumes*, Annales de l’institut Fourier **58** (2008), 447–506.
- [ILTC01] C. Indermitte, Th.M. Liebling, M. Troyanov, and H. ClāfmenÅgon, *Voronoi diagrams on piecewise flat surfaces and an application to biological growth*, Theoretical Computer Science **263** (2001), no. 1, 263 – 274, Combinatorics and Computer Science.
- [JKG07] Miao Jin, Junho Kim, and Xianfeng David Gu, *Discrete surface ricci flow: Theory and applications*, Mathematics of Surfaces XII (Berlin, Heidelberg) (Ralph Martin, Malcolm Sabin, and Joab Winkler, eds.), Springer Berlin Heidelberg, 2007, pp. 209–232.
- [KCD⁺16] Mina Konakovic, Keenan Crane, Bailin Deng, Sofien Bouaziz, Daniel Piker, and Mark Pauly, *Beyond developable: Computational design and fabrication with auxetic materials*, ACM Trans. Graph. **35** (2016), no. 4.
- [KLRR18] Richard Kenyon, Wai Yeung Lam, Sanjay Ramassamy, and Marianna Russkikh, *Dimers and Circle patterns*, arXiv e-prints (2018), arXiv:1810.05616.
- [KPD09] Daniel Kane, Gregory N. Price, and Erik D. Demaine, *A pseudopolynomial algorithm for alexandrov’s theorem*, Algorithms and Data Structures (Berlin, Heidelberg) (Frank Dehne, Marina Gavrilova, Jörg-Rüdiger Sack, and Csaba D. Tóth, eds.), Springer Berlin Heidelberg, 2009, pp. 435–446.
- [KSS06] Liliya Kharevych, Boris Springborn, and Peter Schröder, *Discrete conformal mappings via circle patterns*, ACM Trans. Graph. **25** (2006), no. 2, 412–438.
- [Lis17] Marcin Lis, *Circle patterns and critical Ising models*, arXiv e-prints (2017), arXiv:1712.08736.
- [LP16] Wai Yeung Lam and Ulrich Pinkall, *Holomorphic vector fields and quadratic differentials on planar triangular meshes*, Advances in Discrete Differential Geometry (Alexander I. Bobenko, ed.), Springer Berlin Heidelberg, Berlin, Heidelberg, 2016, pp. 241–265.
- [LP18] ———, *Infinitesimal conformal deformations of triangulated surfaces in space*, Discrete Comput. Geom. **60** (2018), no. 4, 831–858. MR 3869452
- [LPRM02] Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Maillot, *Least Squares Conformal Maps*, ACM Transactions on Graphics **21** (2002), no. 3.
- [Luo04] Feng Luo, *Combinatorial Yamabe flow on surfaces*, Commun. Contemp. Math. **6** (2004), no. 5, 765–780. MR 2100762
- [Luo11] ———, *Rigidity of polyhedral surfaces, iii*, Geom. Topol. **15** (2011), no. 4, 2299–2319.
- [Mer01] C. Mercat, *Discrete Riemann Surfaces and the Ising Model*, Communications in Mathematical Physics **218** (2001), 177–216.

- [MGA⁺17] Haggai Maron, Meirav Galun, Noam Aigerman, Miri Trope, Nadav Dym, Ersin Yumer, Vladimir G. Kim, and Yaron Lipman, *Convolutional neural networks on surfaces via seamless toric covers*, ACM Trans. Graph. **36** (2017), no. 4.
- [Mil94] John Milnor, *The Schläfli Differential Equality*, Collected Papers, vol. 1, Publish or Perish Inc., Houston, 1994, pp. 281–295.
- [MTAD08] Patrick Mullen, Yiyang Tong, Pierre Alliez, and Mathieu Desbrun, *Spectral conformal parameterization*, Proceedings of the Symposium on Geometry Processing (Aire-la-Ville, Switzerland, Switzerland), SGP '08, Eurographics Association, 2008, pp. 1487–1494.
- [Pen87] R. C. Penner, *The Decorated Teichmüller Space of Punctured Surfaces*, Comm. Math. Phys. **113** (1987), no. 2, 299–339.
- [PP93] Ulrich Pinkall and Konrad Polthier, *Computing discrete minimal surfaces and their conjugates*, Experiment. Math. **2** (1993), no. 1, 15–36.
- [Reg61] T. Regge, *General relativity without coordinates*, Il Nuovo Cimento (1955-1965) **19** (1961), no. 3, 558–571.
- [Riv93] Igor Rivin, *On geometry of convex ideal polyhedra in hyperbolic 3-space*, Topology **32** (1993), no. 1, 87 – 92.
- [Riv94] ———, *Euclidean structures on simplicial surfaces and hyperbolic volume*, Annals of Mathematics **139** (1994), no. 3, 553–580.
- [Riv96] ———, *A characterization of ideal polyhedra in hyperbolic 3-space*, Annals of Mathematics **143** (1996), no. 1, 51–70.
- [Roh11] Steffen Rohde, *Oded Schramm: from circle packing to SLE*, Ann. Probab. **39** (2011), no. 5, 1621–1667. MR 2884870
- [RS87] Burt Rodin and Dennis Sullivan, *The Convergence of Circle Packings to the Riemann Mapping*, J. Differential Geom. **26** (1987), no. 2, 349–360.
- [RW84] M. Roček and R. M. Williams, *The quantization of regge calculus*, Zeitschrift für Physik C Particles and Fields **21** (1984), no. 4, 371–381.
- [SC17] Rohan Sawhney and Keenan Crane, *Boundary First Flattening*, ACM Trans. Graph. **36** (2017), no. 5.
- [SC18] Nick Sharp and Keenan Crane, *Variational surface cutting*, ACM Trans. Graph. **37** (2018), no. 4.
- [Sch08] Jean-Marc Schlenker, *Circle patterns on singular surfaces*, Discrete Comput. Geom. **40** (2008), no. 1, 47–102. MR 2429649
- [SLMB05] Alla Sheffer, Bruno Lévy, Maxim Mogilnitsky, and Alexander Bogomyakov, *ABF++ : Fast and Robust Angle Based Flattening*, ACM Trans. Graph. **24** (2005), no. 2.
- [Smi10] Stanislav Smirnov, *Discrete Complex Analysis and Probability*, arXiv e-prints (2010), arXiv:1009.6077.
- [Spr08] Boris A. Springborn, *A variational principle for weighted Delaunay triangulations and hyperideal polyhedra*, J. Differential Geom. **78** (2008), no. 2, 333–367. MR 2394026
- [Spr17] B. Springborn, *Hyperbolic polyhedra and discrete uniformization*, ArXiv e-prints (2017).
- [SS00] A. Sheffer and E. De Sturler, *Surface parameterization for meshing by triangulation flattening*, Proc. 9th International Meshing Roundtable, 2000, pp. 161–172.
- [SSP08] Boris Springborn, Peter Schröder, and Ulrich Pinkall, *Conformal Equivalence of Triangle Meshes*, ACM Trans. Graph. **27** (2008), no. 3.
- [Ste03] Ken Stephenson, *Circle packing: A mathematical tale*, Notices of the American Mathematical Society (2003).
- [Ste05] Kenneth Stephenson, *Introduction to Circle Packing: The Theory of Discrete Analytic Functions*, Cambridge University Press, 2005.
- [SWGL15] Jian Sun, Tianqi Wu, Xianfeng Gu, and Feng Luo, *Discrete conformal deformation: algorithm and experiments*, SIAM J. Imaging Sci. **8** (2015), no. 3, 1421–1456. MR 3369062
- [Thu79] William Thurston, *Geometry and Topology of 3-Manifolds*, Princeton University, 1979.
- [Thu98] William P. Thurston, *Shapes of polyhedra and triangulations of the sphere*, 511–549. MR 1668340
- [Tro91] Marc Troyanov, *Prescribing curvature on compact surfaces with conical singularities*, Transactions of the American Mathematical Society **324** (1991), no. 2, 793–821.

- [VMW15] Amir Vaxman, Christian Müller, and Ofir Weber, *Conformal mesh deformations with möbius transformations*, ACM Trans. Graph. **34** (2015), no. 4, 55:1–55:11.
- [Wil05] S. O. Wilson, *Geometric Structures on the Cochains of a Manifold*, ArXiv Mathematics e-prints (2005).
- [WMKG07] Max Wardetzky, Saurabh Mathur, Felix Kaelberer, and Eitan Grinspun, *Discrete Laplace Operators: No Free Lunch*, Proceedings of the Symposium on Geometry Processing (Alexander Belyaev and Michael Garland, eds.), The Eurographics Association, 2007.
- [WP11] Johannes Wallner and Helmut Pottmann, *Geometric computing for freeform architecture*, Journal of Mathematics in Industry **1** (2011), no. 1, 4.
- [Wu14] Tianqi Wu, *Finiteness of switches in discrete Yamabe flow*, Master’s thesis, Tsinghua University, May 2014.
- [Xu18] Xu Xu, *Rigidity of inversive distance circle packings revisited*, Adv. Math. **332** (2018), 476–509. MR 3810260
- [ZGLG12] Wei Zeng, Ren Guo, Feng Luo, and Xianfeng Gu, *Discrete heat kernel determines discrete riemannian metric*, Graph. Models **74** (2012), no. 4, 121–129.
- [ZGZ⁺14] M. Zhang, R. Guo, W. Zeng, F. Luo, S.-T. Yau, and X. Gu, *The Unified Surface Ricci Flow*, ArXiv e-prints (2014).
- [ZLH⁺18] Hui Zhao, Xuan Li, Ge Huabin, Na Lei, Min Zhang, Xiaoling Wang, and Xianfeng Gu, *Conformal mesh parameterization using discrete calabi flow*, Computer Aided Geometric Design **63** (2018).
- [ZLS07] Rhaleb Zayer, Bruno Lévy, and Hans-Peter Seidel, *Linear Angle Based Parameterization*, Proc. Symp. Geom. Proc., 2007, pp. 135–141.

CARNEGIE MELLON UNIVERSITY, 5000 FORBES AVE, PITTSBURGH, PA 15213