


# Conformal Mesh Deformations with Möbius Transformations

Amir Vaxman  
Vienna University of Technology, Austria

Christian Müller  
Vienna University of Technology, Austria

Ofir Weber  
Bar Ilan University, Israel



**Figure 1:** A panorama of the capabilities of our framework. Deformation of a circular mesh (left). Metric-conformal deformation and interpolation in 3D (center), and an intersection-angle preserving deformation of a planar mesh (right).

## Abstract

We establish a framework to design triangular and circular polygonal meshes by using face-based compatible Möbius transformations. Embracing the viewpoint of surfaces from circles, we characterize discrete conformality for such meshes, in which the invariants are circles, cross-ratios, and mutual intersection angles. Such transformations are important in practice for editing meshes without distortions or loss of details. In addition, they are of substantial theoretical interest in discrete differential geometry. Our framework allows for handle-based deformations, and interpolation between given meshes with controlled conformal error.

**CR Categories:** I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

**Keywords:** Möbius transformations, discrete conformal transformations, circular meshes, handle-based editing, interpolation

## 1 Introduction

Editing discrete surfaces by deformations is an important branch of geometry processing. Meshes are often edited by prescribing a displacement to a subset of vertices, or by interpolating new vertex positions from existing meshes. There are common classes of meshes available for designers, each with a distinctive set of advantages. Triangle meshes are common in computer graphics, and are easy to manipulate. Polygonal meshes are more general and efficient, with less edges and vertices on average, making them popular in industrial and architectural design, where cost effectiveness plays a part. Quadrilateral meshes are viewed as discrete parametrizations of surfaces for their grid structure, and are especially comfortable for subdivision and in computer-aided design. Another important

class is *polyhedral meshes*, where every face is planar. A notable subclass of polyhedral meshes is *circular meshes*, where every face is *concyclic*, i.e., inscribed in a single circle. Considered as discrete curvature-line networks, circular meshes are useful in the field of construction and architectural geometry for their *vertex offset* properties [Liu et al. 2006; Pottmann et al. 2007], and in the field of *discrete differential geometry* for their mathematical properties [Bobenko and Suris 2008]. Circular meshes allow for proper discrete definitions of Gauss maps, shape operators, and conformal transformations. Note that triangle meshes are circular meshes by definition. Our framework is inspired by the viewpoint of *surfaces from circles* [Bobenko et al. 2006], and we show its usefulness in both design and theory. For this purpose, we employ notions from *Möbius geometry*, in which circles and spheres are fundamental objects.

A common method for mesh editing is by positional handles: the user picks and drags a small set of vertices, and the other vertices transform respecting fairness measures and constraints. Another method is the creation of meshes as averages of existing shapes by interpolation. We expect the interpolated shapes to conform to the same constraints as the originals. Moreover, if constraints are violated in the original shapes, we expect this violation to be bounded in the interpolated shapes, for a stable and intuitive result.

A common quality measure is *conformality*. Conformal transformations comprise local similarities, and consequently preserve fine details and texture (see Figure 2 for a demonstration). However, their main disadvantage is the introduction of scale variations, which might produce unintuitive results in practice. We offer a way to control and balance such variations. Conformal transformations of continuous surfaces are well defined, but there are several approaches to what discrete conformality constitutes. Moreover, some approaches are only defined in two dimensions, and cannot be generalized to 3D. We offer a unified and practical approach to several of these definitions, for both two and three dimensions.

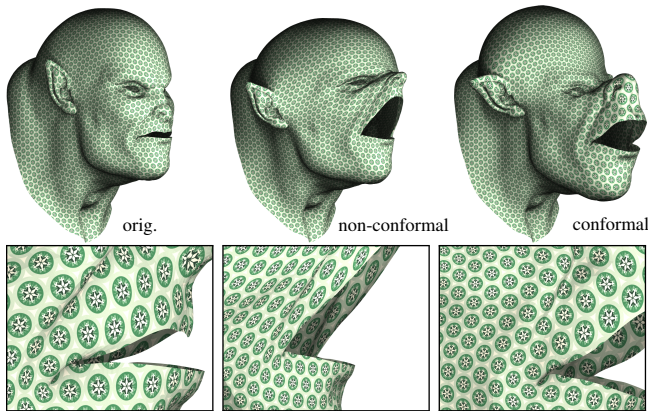
Finally, editing circular meshes is difficult, as face concyclicity depends nonlinearly on the vertex positions. Two approaches handle this difficulty: projection of generally-deformed meshes into nearby circular meshes, and editing circular meshes in a designated subspace that preserves concyclicity. We take the latter approach.

**Our contributions** We offer an editing framework for triangle and circular polygonal meshes by using face-based piecewise-compatible Möbius transformations in 2D and 3D. Our approach allows for:

1. A definition and optimization of discrete conformality that unifies the following approaches:

- Preservation of circular meshes. [Bobenko et al. 2006].

- Discrete-conformal equivalence by metric scaling of edge lengths [Springborn et al. 2008].
  - Preservation of the angles between face circumcircles, as a part of the circle pattern approach [Bobenko and Springborn 2004; Kharevych et al. 2006].
2. Mesh editing by prescribing positional constraints, preserving facewise concyclicity.
  3. Interpolation and extrapolation between two shapes, both concyclicity preserving and with bounded discrete-conformal distortion.
  4. A simple Gauss-Newton based optimization that fits all tasks.



**Figure 2:** The beast head is deformed conformally and non-conformally with the same positional constraints. The conformal deformation introduces some scale but avoids shear and preserves texture details more faithfully.

## 2 Related Work

**Discrete conformal mesh transformations** There are several definitions of discrete conformality, each introducing discrete counterparts to continuous conformal invariants. It is common to consider the singular values of the piecewise constant face-based Jacobian of a mapping between two meshes. The ratio of the maximal to the minimal singular values is denoted as the *quasiconformal error*, where 1 indicates a perfect similarity transformation. Several methods attempt to optimize for face-based isometry [Alexa et al. 2000] or similarity [Liu et al. 2008] by alternating between a local projection of a transformed element onto the closest similar or rigid counterpart, followed by a global reconstruction for the vertices of the mesh by a least-squares approximation. Alternatively, the mesh is encoded as a set of local frames [Lipman et al. 2007; Paries et al. 2007]. Several recent methods target mappings that are guaranteed to be locally injective and with low distortion [Lipman 2012; Weber et al. 2012; Schüller et al. 2013; Weber and Zorin 2014]. However, in contrast to our method, these methods are not designed to compute a mapping  $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ . Our method deals with planar, as well as three dimensional mappings in a unified framework. [Crane et al. 2011] compute conformal deformations of surfaces embedded in  $\mathbb{R}^3$  using quaternionic spin transformations, requiring only a linear solve. However, it is impossible to prescribe positional constraints.

A discrete theory of conformal transformations and invariants is explored in the field of discrete differential geometry [Bobenko and Springborn 2004; Kharevych et al. 2006], where surfaces are represented as *circle patterns*. We employ such an approach in this work. The adjacency relations and the mutual intersections of circle patterns are conformal invariants in Möbius geometry. Applications of the circle-pattern approach include parametrization [Kharevych

et al. 2006], and constructing surfaces from spherical meshes by duality [Bobenko et al. 2006; Müller 2011]. A primary invariant is the quadrilateral *cross-ratio*. Quadrilaterals with similar cross-ratios are thus considered conformally equivalent. Specifically, *conformal quads* are conformally equivalent to regular quads, and are the basic blocks for discrete isothermic surfaces [Bobenko and Suris 2008]. Similar approaches are applied for hexagonal meshes [Müller 2011], where conformal hexagons are defined as pairs of conformal quads. Discrete conformal invariants are preserved under global transformations in such discretizations, but these methods do not offer a flexible approach to mesh editing. A related approach to discrete conformality in triangle meshes rescales the discrete metric (edge lengths). It is used to achieve a conformally-equivalent flat or constant-curvature metric, for uniformization and parametrization of meshes [Gu and Yau 2002; Ben-Chen et al. 2008; Springborn et al. 2008]. However, metric scaling by itself is not enough to define a conformal deformation of surfaces in  $\mathbb{R}^3$ , and these methods do not support prescription of positional constraints. We offer a unification of the circle-pattern and the metric-scaling approaches by intuitive reductions of piecewise-Möbius transformations.

**Editing circular meshes** There are two main approaches to the design of circular meshes. One approach exploits *parallel meshes* of a given mesh. Parallel meshes are a rather restricted approach to mesh editing, albeit preserving concyclicity and maintaining other important discrete notions of curvature and duality. Methods that specifically pertain to polyhedral meshes, such as using face-based projective maps [Vaxman 2014], cannot address circular meshes since projective transformations do not preserve circles. Another approach addresses general deformations of constrained meshes, either by first and second-order approximations [Habbecke and Kobbelt 2012; Yang et al. 2011], which only produce approximately-constrained meshes, or by means of projection: deform a mesh freely into a new shape, and then find a similar shape that satisfies constraints, such as face concyclicity. ShapeUp [Bouaziz et al. 2012] alternates between global integration and local element optimizations, while [Tang et al. 2014] solves a Gauss-Newton-type system, adding auxiliary variables to make all constraints quadratic. Both methods present a way to optimize a given deformation into a circle-preserving one, but do not attempt to either characterize the class of circle-preserving transformations, or derive meaningful discrete conformal variations of such deformations. Our Möbius-based method parametrizes the space of circle and cross-ratio preserving transformations, providing an efficient optimization for exploring this space while imposing positional constraints.

**Interpolation methods** Interpolated meshes are typically computed by the extraction of infinitesimal geometric quantities from two (or more) meshes, blending them, and then reconstructing vertex positions. As-rigid-as-possible interpolation [Alexa et al. 2000] factors and interpolates the facewise affine mappings between corresponding triangles with a polar decomposition. Interpolating large rotations is problematic due to ambiguities in the choice of rotation angles. Shape-space based methods, such as [Kilian et al. 2007], navigate the space of possible deformations between two shapes, by finding geodesics of a deformation metric. However, such methods are computationally expensive, due to an intricate nonlinear optimization. Recent methods avoid the rotational interpolation problems by interpolating *intrinsic* geometric quantities, such as lengths in two dimensions [Chen et al. 2013], and in addition dihedral angles in three dimensions [Winkler et al. 2010]. Reconstruction with pointwise bounded conformal distortion is provided by [Chen et al. 2013], but the method is not applicable to three dimensions. Our method avoids dealing with rotation ambiguities by blending cross-ratio differences and is applicable in two and three dimensions.

### 3 Complex Möbius Transformations

We introduce our framework with piecewise-Möbius transformations, parametrizing the shape space of admissible transformations. We begin with Möbius transformations in the complex plane.

#### 3.1 Transformations of single polygons

Let  $\mathcal{M} = \{\mathcal{V}, \mathcal{E}, \mathcal{F}\}$  be a mesh of arbitrary topology and face degrees (i.e., number of vertices in every face) embedded in the complex plane, such that the position of every vertex  $v_i \in \mathcal{V}$  is a complex number  $z_i \in \mathbb{C}$ . We consider deformations that are characterized by a single Möbius transformation per face. A Möbius transformation in the complex plane is represented by a *linear-fractional transformation*  $m_f : \hat{\mathbb{C}} \rightarrow \hat{\mathbb{C}}$ , ( $\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ ) parametrized by coefficients  $\{a_f, b_f, c_f, d_f\} \in \mathbb{C}$  s.t.  $a_f d_f - b_f c_f \neq 0$ . The positions  $z_i \in f$  of a single face  $f \in \mathcal{F}$  are transformed by:

$$m_f(z_i) = (a_f z_i + b_f)(c_f z_i + d_f)^{-1}.$$

It is important to note that we define the transformation as applying only to the *vertices* in the face, rather than to its edges (which in general, do not transform into straight lines). We thus consider the transformed edges as the straight segments between the transformed vertices. The edge transformations have a simple closed-form expression, which is the basis for our framework. For all  $(z_i, z_k) \in f$ :

$$m_f(z_k) - m_f(z_i) = \frac{a_f d_f - b_f c_f}{(c_f z_i + d_f)(c_f z_k + d_f)} (z_k - z_i).$$

The group of Möbius transformations contains all similarities and inversions in spheres. Specifically, a Möbius transformation is a similarity if and only if  $c_f = 0$ .

**Degrees of freedom** The coefficients represent a homogeneous transformation, defined up to a global complex multiplicative factor. Assuming the transformation is nonsingular, we can *normalize* the transformation by setting  $a_f d_f - b_f c_f = 1$ , which results in three complex (six real) degrees of freedom. Hence, a Möbius transformation is uniquely defined by repositioning three vertices.

**Corner reciprocals** We use the common term ‘‘corner’’ for the pair consisting of a vertex and an adjacent face. Throughout the paper we denote the reciprocal to the denominator of a normalized Möbius transformation as:  $X_{f,i} = (c_f z_i + d_f)^{-1}$ . With this definition, the transformation of an edge can be concisely written as:

$$m_f(z_k) - m_f(z_i) = z_{ik} X_{f,i} X_{f,k}, \quad (1)$$


where  $z_{ik} = z_k - z_i$ . Interestingly, the edge-vector transformation depends only on the denominator, while the numerator encodes the translation. That means that the full transformation can be determined from the denominator and the repositioning of one vertex. The corner variables encode local scaling and rotation: the edge  $z_{ik}$  is scaled by  $|X_{f,i} X_{f,k}|$ , and rotated by  $\arg(X_{f,i} X_{f,k})$ . We make use of this quantification of rotation and scaling for our definition of discrete conformality in the next section.

**Signs** The coefficients of a normalized Möbius transformation are only unique up to a multiplication with  $-1$ . There is no difference with regard to the transformations of either vertices or edges, but the corner reciprocals  $X = (cz + d)^{-1}$  in a single face  $f$  are unique only up to sign. We explain how to deal with this ambiguity later.

#### 3.2 Piecewise-compatible transformations

Consider two adjacent faces  $f, g$  with a common edge  $ik$  (see Figure 3), such that their respective common vertices  $z_i, z_k$  are transformed via two normalized Möbius transformations  $m_f$  and  $m_g$ . In order for the two transformations to be *compatible*,  $z_i$  and  $z_k$  must transform to the same target positions by both transformations. A necessary condition for this compatibility is that the common edge vector is transformed to the same target vector on both sides:


$$w_k - w_i = z_{ik}(X_{f,i} X_{f,k}) = z_{ik}(X_{g,i} X_{g,k}). \quad (2)$$



**Figure 3:** A triangle flap (left) with the notation we use throughout. Our claims are made with polygonal meshes (right) by choosing arbitrary opposite points  $j, l$ , which cancel out in our derivations.

If the transformed edge vectors are explicitly represented by vertex  $w$  values, the condition above is also sufficient (the edge vectors constitute an *exact 1-form*). The original edges  $z_{ik}$  may seem redundant in the formulation, if  $w_{ik}$  are not considered, but we always keep them for clarity since they are not reducible in three-dimensions. We denote a set of corner variables  $X$  which obeys Equation (2) as *integrable*. Such a set of  $X$  variables uniquely defines a deformed mesh up to a global translation.

A transformation of a polygonal mesh where every face is related to its image by a Möbius transformation is denoted as a *piecewise-compatible Möbius* (abbr. **PCM**) transformation. See Figure 4 for an example of a PCM transformation of a quad mesh. Every deformation of triangle meshes in  $\mathbb{C}$  is a PCM transformation by definition, and it may seem unintuitive to parametrize triangle mesh deformations in this nonlinear manner. In the next section, we show that PCM transformations are valuable for all types of meshes, as they allow for elegant definitions of discrete conformality, which can then be optimized for with ease.



**Figure 4:** Quad meshes related by a piecewise-compatible Möbius transformation.

## 4 Discrete Conformality

A guiding principle for terming transformations as discrete conformal is the identification of discrete quantities which are considered as conformal invariants, and which such transformations preserve. Our fundamental invariant for definitions of conformality is the *cross-ratio* of vertices.

#### 4.1 The complex cross-ratio


A Möbius transformation is synonymous with the preservation of the complex *cross-ratio*, which consequently results in the preservation



of circles. Given points  $z_i, z_j, z_k, z_l$  transformed by a Möbius transformation to respective points  $w$ , we get that:

$$\text{cr}_z[i, j, k, l] = \frac{(z_i - z_j)(z_k - z_l)}{(z_j - z_k)(z_l - z_i)} = \text{cr}_w[i, j, k, l]. \quad (3)$$

Therefore, two polygonal faces of more than three vertices are Möbius-equivalent if and only if every four corresponding vertices have the same cross-ratio. For instance, a regular square has a real cross-ratio of  $-1$ , and all equivalent quads are sometimes called “conformal squares” [Müller 2011] (see Figure 5).



**Figure 5:** A regular quad (left) and its conformally equivalent (center). Positive cross-ratios mark self-intersecting quads (right).

**Cross-ratio values** A quad is concyclic if and only if its complex cross-ratio is real. Moreover, the cross-ratio is negative if and only if the four vertices are arranged on their circumcircle in cyclic order (clockwise or counter clockwise). That means that self-intersecting concyclic quads are identified with positive cross-ratios.

**Under compatible Möbius transformations** Consider the notation of Figure 3, and a PCM transformation parametrized by corner variables  $X$ . Plugging Equation (1) in the cross-ratio definition, we get:  $\text{cr}_w[i, j, k, l] = \text{cr}_z[i, j, k, l] (X_{g,k} X_{f,i})^{-1} X_{g,i} X_{f,k}$ . Compatibility conditions (2) state that:

$$X_{f,i} X_{f,k} = X_{g,i} X_{g,k} \text{ or } X_{f,i}/X_{g,i} = X_{g,k}/X_{f,k}. \quad (4)$$

And thus  $\text{cr}_w[i, j, k, l]$  equals:

$$\text{cr}_z[i, j, k, l] (X_{f,k}/X_{g,k})^2 = \text{cr}_z[i, j, k, l] (X_{g,i}/X_{f,i})^2. \quad (5)$$

We get an interesting relation: the change in the cross-ratio based on the edge  $ik$  solely depends on the quotient of corner reciprocals neighboring at either vertex  $i$  or  $k$  (assuming the transformations are compatible). This relation is symmetric on both sides of the edge.

Together with Equation (2) we get that the two faces  $f, g$  are transformed by a single Möbius transformation if and only if, in vector notation,  $(X_{g,k}, X_{g,i}) = \pm(X_{f,k}, X_{f,i})$  and thus  $(c_f, d_f) = \pm(c_g, d_g)$ . Remember that the corner variables in a face are invariant to a uniform change of sign, and thus we can safely say that the transformations of faces  $f, g$  are identical if and only if there are corner variables such that  $(X_{g,k}, X_{g,i}) = (X_{f,k}, X_{f,i})$ .


## 4.2 Combining notions of discrete conformality

The cross-ratio is a measure of Möbius equivalence that also encodes common notions in discrete conformal geometry: the *length cross-ratio* [Springborn et al. 2008], and the *circle-intersection angle* [Bobenko and Springborn 2004].

**Length cross-ratio** The *modulus* of the cross-ratio is simply the ratio of the respective edge lengths  $\text{lcr}_z[i, j, k, l] = |\text{cr}_z[i, j, k, l]|$ . Two combinatorially-equivalent triangle meshes that preserve the length cross-ratio of the quad around every inner edge are denoted as “discrete-conformally equivalent” by [Springborn et al. 2008]. In order to avoid notational confusion, we denote such surface transformations as *metric conformal* (or **MC** in short). In the notation of Figure 3, we only preserve  $\text{lcr}_z[i, j, k, l]$  (and its equivalent permutations), and not, e.g.  $\text{lcr}_z[i, k, j, l]$ . Möbius transformations are trivially MC transformations.

**Conformal factors** An equivalent way to define MC transformations [Springborn et al. 2008] is by parametrizing them with *conformal factors*. Two meshes  $\mathcal{M}_1, \mathcal{M}_2$  are MC equivalent if and only if there is a single scalar value  $u_i$  for every vertex  $z_i$ , denoted the conformal factor, such that the length of each edge transforms as  $|w_{ij}| = |z_{ij}| \exp((u_i + u_j)/2)$ .

**MC conformality in our framework** From Equation (5) and the preservation of the length cross-ratio, we can deduce that a piecewise-compatible Möbius transformation is also MC if and only if for every two adjacent corner reciprocals  $X_{f,i}, X_{g,i}$ , we have that  $|X_{f,i}|^2 = |X_{g,i}|^2$ . Furthermore, the conformal factor  $u_i$  equals  $2 \log(|X_{f,i}|)$  (also for all adjacent corners), from the transformation of edges in Equation (1). Thus, MC transformations are an intuitive subclass of PCM transformations where for every vertex, all adjacent corner reciprocals have the same modulus.



**Figure 6:** Two adjacent triangles  $f = ijk, g = kli$  and the intersection angle  $\phi_{ik}$  between their respective circumcircles.

**Phase cross-ratio** In a complementary manner, the phase of the cross-ratio is established as [Bobenko and Schröder 2005]:

$$\text{cr}_z[i, j, k, l] = \text{lcr}_z \exp(i(\pi - \phi_{ik})), \text{ or } \cos(\phi_{ik}) = -\frac{\text{Re}(\text{cr}_z)}{\text{lcr}_z},$$

where  $\phi_{ik}$  is the *intersection angle* of the two circumcircles defined by  $ijk$  and  $kli$  (see Figure 6), preserved by Möbius transformations as well. PCM transformations preserve  $\phi_{ik}$  and only if  $X_{f,i} X_{g,i}^{-1} \in \mathbb{R}$  (having the same phase), and thus  $\text{cr}_w[i, j, k, l] (\text{cr}_z[i, j, k, l])^{-1} \in \mathbb{R}$ . We denote such a transformation as *intersection-angle preserving* (in short, **IAP**).

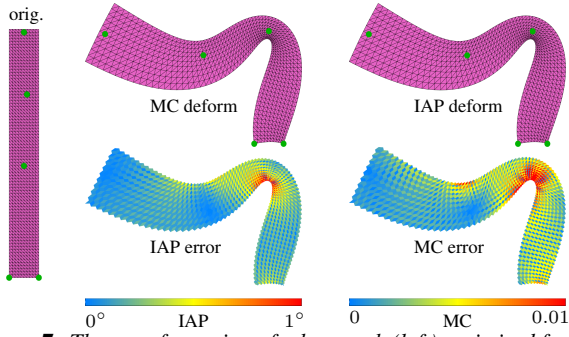
To conclude, we see that the PCM subclasses of MC (preserving cross-ratio modulus) and IAP (preserving phase) transformations are complementary definitions of discrete conformality. If a transformation is both IAP and MC, then all corner variables adjacent at a vertex are equal, all complex cross-ratios are preserved, and thus the entire transformation is a *single* Möbius transformation. Although this connection between MC and IAP can be deduced from the cross-ratio, we are not aware of any previous attempt to quantify it in deformations with corner reciprocals in the way we do. We demonstrate deformations that preserve MC and IAP (separately) in Figure 7. These properties are also surprisingly preserved in three dimensions with only a few differences, as we show next.

## 5 Quaternionic Möbius transformations

We generalize piecewise-Möbius transformations and discrete conformality to three-dimensions with *quaternionic* variables and transformations. Quaternions  $q = (r, v) \in \mathbb{H}$  are defined by a scalar part  $r \in \mathbb{R}$  and a vector part  $v \in \mathbb{R}^3$ , parametrizing  $\mathbb{R}^4$ . We identify  $\mathbb{R}^3$  with the space of *imaginary* quaternions  $\text{Im } \mathbb{H} = \{(r, v) \in \mathbb{H} \mid r = 0\}$ . Hence, quaternionic transformations in three-dimensions can be understood as restrictions of mappings  $\mathbb{H} \rightarrow \mathbb{H}$  to the space of imaginary quaternions  $\text{Im } \mathbb{H}$ , thus called *imaginary preserving*.

In the polar representation, quaternions are  $s(\cos(\phi), v \sin(\phi))$ , and are defined by a positive modulus  $s$ , a phase  $\phi$ , and a unit direction  $v \in \mathbb{R}^3$ . The two interesting quantities for discrete conformality





**Figure 7:** The transformation of a bar mesh (left) optimized for MC (middle) and IAP (right). We show the respective transformations with (green) positional constraints (upper images), IAP error (lower middle), measured as  $|\phi_{\text{new}} - \phi_{\text{old}}|$  on every edge, and MC error (lower right), measured as  $|\text{lcr}_w / \text{lcr}_z - 1|$ . Error images only show internal edges. We show IAP error only for the MC-optimized mesh and vice versa, since the optimized measures are below  $10^{-7}$ .

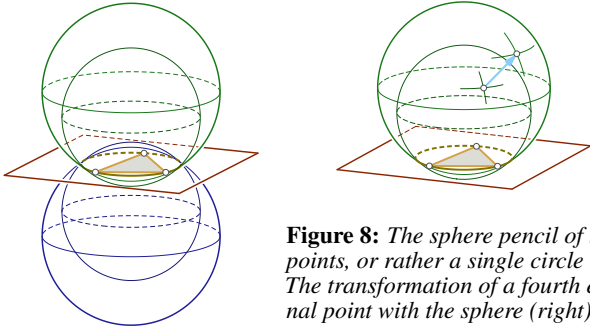
are the modulus and the phase, corresponding directly to respective properties in the complex plane. Fortunately, the transition from complex numbers to quaternions is mostly straightforward.

### 5.1 Möbius transformations in quaternionic variables

Quaternionic Möbius transformations can be parametrized by quaternionic coefficients  $a, b, c, d \in \mathbb{H}$  in several equivalent forms, from which we choose the following:

$$m(q) = (aq + b)(cq + d)^{-1}.$$

This definition is used as a quaternionic counterpart to the complex Möbius transformation in several works [Wilker 1993; Hertrich-Jeromin 2003; Jakobs and Krieg 2010]. The coefficients are unique up to a multiplication with a real scalar (but not with a quaternion, due to noncommutativity), and thus there are 15 degrees of freedom.



**Figure 8:** The sphere pencil of three points, or rather a single circle (left). The transformation of a fourth external point with the sphere (right).

**Imaginary-preserving transformations** Transformation  $m$  maps  $\mathbb{H} \rightarrow \mathbb{H}$ , while we require that  $m(q) \in \text{Im } \mathbb{H}$  for all  $q \in \text{Im } \mathbb{H}$ . This leads to the following necessary and sufficient quadratic conditions [Jakobs and Krieg 2010]:

$$a\bar{c} \in \text{Im } \mathbb{H}, \quad b\bar{d} \in \text{Im } \mathbb{H}, \quad \bar{a}d - \bar{b}c \in \mathbb{R}. \quad (6)$$

We include a proof in Appendix A. These 5 real conditions leave 10 degrees of freedom for imaginary-preserving transformations. The repositioning of three points in  $\mathbb{R}^3$  claims 9 degrees, leaving out a single degree of freedom, which is not enough for the repositioning of an extra fourth point. Thus, unlike the complex case, repositioning a triangle does not uniquely define the Möbius transformation, while a fourth point overconstrains it. In order to understand the geometric meaning of these DOF in  $\mathbb{R}^3$ , note that Möbius transformations in

$\mathbb{R}^3$  transform generalized 2-spheres into others (2-hyperplanes are generalized 2-spheres containing the point in infinity). Without loss of generality, consider Möbius transformations which fix three given points in  $\mathbb{R}^3$ . These points define a one-parameter set of spheres (also known as a *sphere pencil*, see Figure 8). The transformations under considerations have a degree of freedom in transforming a sphere  $S_1$  in the pencil into another sphere  $S_2$ , but they do not allow the prescription of an arbitrary point on  $S_1$  into  $S_2$ . Consequently, there is a one-parameter family of transformations that map a concyclic face into another. We deal with this ambiguity further on.

**Edge compatibility conditions** For an imaginary-preserving Möbius transformation with  $\delta = \bar{a}d - \bar{b}c \in \mathbb{R}$  we have:

$$\begin{aligned} m(q_k) - m(q_i) &= \delta^{-1}(\overline{cq_i + d})^{-1}(q_k - q_i)(cq_k + d)^{-1} \\ &= \delta^{-1}(\overline{cq_k + d})^{-1}(q_k - q_i)(cq_i + d)^{-1}, \end{aligned}$$


which directly generalizes our result in the complex plane (Equation (2)). For a complete proof see Appendix A. Note that this equation is surprisingly symmetric with relation to the positions of  $cq_k + d$  and  $cq_i + d$ . That is because  $m(q_k) - m(q_i) \in \text{Im } \mathbb{H}$ , and thus  $m(q_i) - m(q_k) = m(q_k) - m(q_i)$ .

**Discussion**  $\delta = \bar{a}d - \bar{b}c$  is real by condition (6). Since the coefficients are unique up to a real multiplier, we can only normalize transformations to have  $\delta = \pm 1$ . In fact, the sign of  $\delta$  splits Möbius transformations into two disjoint sets. Consider a transformation with  $\delta = -1$  acting on a circle. It is possible to construct another transformation with  $\delta = 1$  by reflection in the plane containing the circle, leaving it fixed. As we only care about the effect of Möbius transformations on the concyclic faces, we only assume to work on  $\delta = 1$  transformations as a comfortable choice, and optimize accordingly; this does not limit our degrees of freedom.

Analogous to the complex case, and in the notations of Figure 3, we define quaternionic corner reciprocals  $X_{f,i|k} = (c_f q_{i|k} + d_f)^{-1}$  for the positions  $q_i, q_k$  of adjacent vertices of a face  $f$  such that:

$$m(q_k) - m(q_i) = \bar{X}_{f,i} q_{ik} X_{f,k} = \bar{X}_{f,k} q_{ik} X_{f,i}. \quad (7)$$

These are our edge-compatibility conditions for PCM transformations. Naturally, if  $X_{f,i} = X_{g,i}$  and resp. for vertex  $k$ , both faces are related by a single Möbius transformation. Since there is a one-parameter choice of transformations per face, the converse is not strictly-speaking correct. However, this is irrelevant practically, as we employ the reciprocals directly, and thus, e.g., demand  $X_{f,i} = X_{g,i}$  in our computations if we need to, as done in two dimensions. We depict quaternionic PCM transformations in Figure 9.



**Figure 9:** Eye mesh (left) and PCM deformation (right). The relative concyclicity of faces is preserved.

## 5.2 Quaternionic cross-ratio and discrete conformality

Given indices  $i, j, k, l$ , the quaternionic cross-ratio can be defined as [Hertrich-Jeromin 2003, Sec. 4.9]:

$$\text{cr}_q[i, j, k, l] = (q_i - q_j)(q_j - q_k)^{-1}(q_k - q_l)(q_l - q_i)^{-1}.$$

The quaternionic cross-ratio is not fully preserved under a Möbius transformation, but is rather conjugated by  $(cq_i + d)$  upon the transformation from  $q$  to  $w$ :

$$\text{cr}_w[i, j, k, l] = (cq_i + d) \text{cr}_q[i, j, k, l] (cq_i + d)^{-1}.$$

The converse is also true: if a cross-ratio is conjugated by some quaternion, there is a Möbius transformation that induces  $q \rightarrow w$ . It is evident that the length cross-ratio (modulus) and the phase cross-ratio, calculated as  $\cos(\phi_{ik}) = -\text{Re}(\text{cr})/|\text{cr}|$  [Bobenko and Schröder 2005] (see Figure 6), are preserved.

The cross-ratio  $\text{cr}_w[i, j, k, l]$ , under a PCM transformation, cannot be decoupled into corner reciprocals  $X$  and the original cross-ratio  $\text{cr}_q$  as easily as the complex one, due to the noncommutativity of multiplication. However, we can form simple expressions for its invariants with certain conditions:

*MC Transformations:* The length cross-ratio is preserved if there exists a PCM transformation with corner variables  $X$  s.t.  $|X_{f,i}| = |X_{g,i}|$  and resp. for  $k$ , exactly like in the complex plane.

*IAP transformations:* The phase cross-ratio is preserved when the new cross-ratio is a similarity of the old one (i.e.  $\text{cr}_w = u \text{cr}_q \bar{u}$  for some  $u \in \mathbb{H}$ ). Thus, if there exists a PCM transformation s.t.  $X_{g,i} X_{f,i}^{-1} \in \mathbb{R}$  and resp. for  $k$ , then it is also IAP.

IAP and MC are complementary in quaternions as well. However, note that the IAP condition is much stronger in  $\mathbb{R}^3$ ; the two corner variables have to be the same up to modulus, whereas in the MC condition, these variables only share a single modulus. In fact, IAP transformations are then only at most one scalar degree of freedom per edge weaker than a *single* Möbius transformation of the entire mesh, and are therefore difficult to attain. This makes IAP conformality too strong in three-dimensions, and we do not use it.

## 6 Editing with Möbius Transformations

Having defined compatibility and conformality, we next utilize them for our applications: editing triangular and circular polygonal meshes conformally with positional constraints, and interpolation between meshes with controlled conformal distortion. We minimize energies with linear and quadratic equality constraints at most; the optimization of which is shown to be effective in [Tang et al. 2014]. We explain the generic process of optimization we utilize for all our operations in Section 8. For clarity of reading, we include a quick reference guide in the supplementary material.

### 6.1 Setting

Our input is the original positions of the vertices ( $z$  in 2D and  $q$  in 3D), and a set of indices of positional constraints  $w^*$ . We output the deformed vertex positions  $w$  (complex or quaternionic, according to context). In our framework we use two systems of variables, each with specific advantages for editing purposes:

**Corner Variables** In this system, we explicitly use the corner reciprocals  $X$  in addition to the positions  $w$  as variables. Therefore, we need to constrain the edge compatibility conditions (2) (complex) or (7) (quaternionic). We call this setting of variables *corner variable minimization* (abbr. **CV min.**). In a triangle mesh we have  $3|\mathcal{F}|$  corner variables and  $|\mathcal{V}|$  positional variables summing at  $3|\mathcal{F}| + |\mathcal{V}| \approx 7|\mathcal{V}|$  in a semi-regular mesh (vertices of mostly valence 6).

**Edge Deviation** Instead of corner reciprocals, we may utilize *vertex-based reciprocals*  $Y_v, \forall v \in \mathcal{V}$ , and use the edge compatibility conditions as soft constraints. We quantify the exact transformation of the edges by using multiplicative *edge deviations*  $e_{ik}$  such that:

$$e_{ik} w_{ik} = Y_i z_{ik} Y_k \text{ in 2D and } e_{ik} w_{ik} = \bar{Y}_i q_{ik} Y_k \text{ in 3D.}$$

In our basic optimization, only the  $Y$  and  $w$  variables are needed, counting to only  $2|\mathcal{V}|$  variables for triangle meshes. With edge deviations, the variables count to  $2|\mathcal{V}| + |\mathcal{E}| \approx 5|\mathcal{V}|$ . We denote this system as *edge deviation minimization* (abbr. **ED min.**). Thus, we use the ED system mainly to save variables.


### 6.2 As-Möbius-as-possible deformations

We consider a transformation that preserves both length and phase cross ratios (MC + IAP) as an ideal discrete conformal transformation, equivalent to applying a Möbius transformation to a 1-ring neighborhood of a vertex. This is the case when all the adjacent corner variables are equal. In order to construct transformations that are ideally Möbius, we formulate the *as-Möbius-as-possible* (**AMAP**) energy. We state the minimization problem in both CV and ED systems as follows ( $w_{ik}$  is short for the explicit term  $w_k - w_i$ ):

	$\forall f, g \in \mathcal{F}, f, g \in N(i)$	s.t. $\forall ik \in (f, g)$
2D	$E_{\text{AMAP,CV}} = \sum  X_{f,i} - X_{g,i} ^2$	$w_{ik} = X_{f,i} z_{ik} X_{g,k}$
3D		$w_{ik} = \bar{X}_{f,i} q_{ik} X_{g,k}$
2D	$E_{\text{AMAP,ED}} =  w_{ik} - Y_i z_{ik} Y_k ^2$	Unconstrained
3D		

Note that in this basic AMAP minimization, the ED system is superior to the CV system since minimization is unconstrained, and we use the minimal amount of variables (no edge deviations needed).

**Positional constraints** We use the explicit positions  $w$ , and have linear constraints:  $C_{\text{POS}}(i) = w_i - w_i^* = 0$  for all handles. We do so instead of variable substitution, since we begin minimizing from a previous solution that does not obey these constraints.

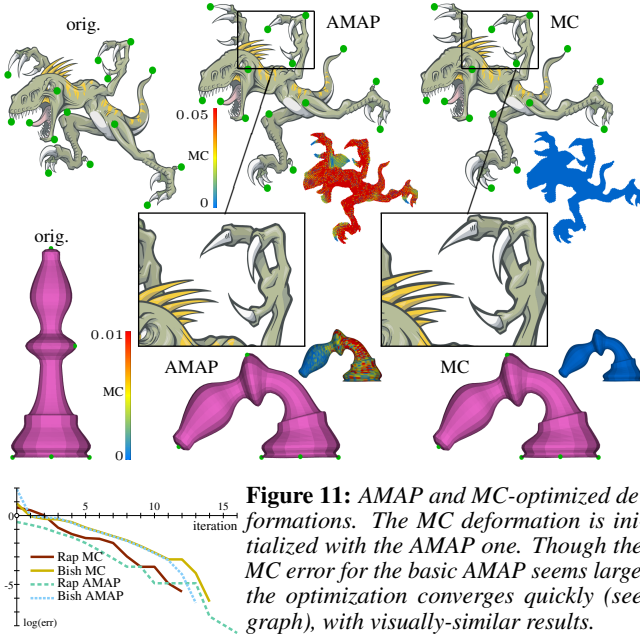


**Figure 10:** AMAP deformation in 2D. Original mesh (left), basic AMAP deformation (center) and inversion weighting with  $\alpha_{\text{INV}} = 0.1$  (right). The middle mesh introduces a large Möbius inversion in the arms, resulting in a quasiconformal error. Inversion control homogenizes the scale, at the cost of less metric conformality.

**Inversion weighting** AMAP deformations strive to reduce the differences between the piecewise transformations, in order to reduce discrete conformal energy. However, adherence to conformality often introduces undesirable scale variations (see Figure 10). A user should practically be able to balance uniform scaling against strict conformality. Scale variations are associated with inversions in the transformations, which are avoided if the Möbius coefficient  $c_f = 0$ . This is equivalent to when all corner reciprocals (vertex reciprocals in ED minimization) of adjacent vertices in a face are equal. Thus, to reduce inversions, we add an energy term for each edge:

$$E_{\text{INV}} = \alpha_{\text{INV}} \sum_{(i,k) \in \mathcal{E}} |Y_i - Y_k|^2,$$

and similarly for all two neighboring corner variables in the CV system.  $\alpha_{\text{INV}}$  weighs the amount of inversion we would like to



**Figure 11: AMAP and MC-optimized deformations.** The MC deformation is initialized with the AMAP one. Though the MC error for the basic AMAP seems large, the optimization converges quickly (see graph), with visually-similar results.

reduce. Unless otherwise stated, we always use  $\alpha_{INV} = 0.1$  for 2D, and  $\alpha_{INV} = 0.5$  for 3D. See Figure 10 for an example.

### 6.3 MC, IAP, and polygonal circular meshes

**Corner-variable minimization** MC optimization is done by constraining  $|X_{f,i}|^2 = |X_{g,i}|^2$ . Optimizing for IAP requires the constraint  $X_{g,i}^{-1}X_{f,i} \in \mathbb{R}$ . The latter is not quadratic, but we can substitute it with  $\bar{X}_{g,i}X_{f,i} \in \mathbb{R}$ , since  $q^{-1} = \bar{q}/|q|^2$ . Polygonal faces transform with a single Möbius transformation by default, since we use explicit corner variables within each face.

**Edge-deviation minimization** We need to use the aforementioned edge deviations  $e_{ik}$  explicitly. The complex cross-ratio resulting from the transformation is:

$$\text{cr}_w[i, j, k, l] = \text{cr}_z[i, j, k, l] e_{jk} e_{li} (e_{ij} e_{kl})^{-1}.$$

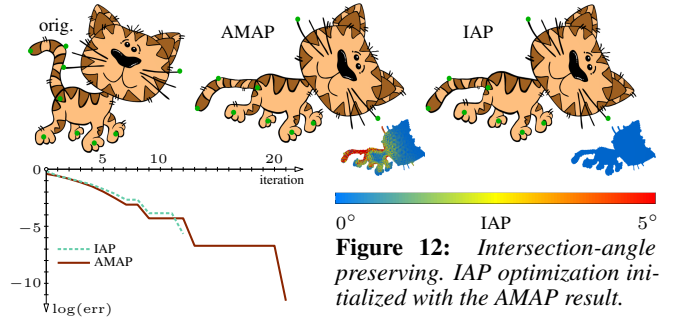
Constraining  $|e_{ik}|^2 = 1$  for each edge  $ik$  results in a MC transformation (preserving the modulus), and constraining  $e_{ik} \in \mathbb{R}$  results in an IAP transformation. Constraining  $e_{ij} e_{kl} = e_{jk} e_{li}$  preserves the full cross-ratio, for a single Möbius transformation of  $ijkl$ . We then constrain it for every four consecutive vertices in every polygon.

However, the quaternionic cross-ratio does not transform similarly. It is straightforward to check that the IAP and MC conditions are equivalent, but the condition for full cross-ratios is not quadratic. Thus, the ED system is suboptimal for polygonal meshes in 3D.

In our examples, and for simplicity of exposition, we use the ED system for all 2D deformations, and the CV system for 3D deformations. We found this balance to be the easiest to utilize in an implementation. We exemplify MC deformations in Figure 11, IAP in Figure 12, and polygonal circular meshes in Figure 13.

## 7 Interpolation

We work with two corresponding meshes  $\mathcal{M}_0, \mathcal{M}_1$ , at times  $t = 0$  and  $t = 1$  respectively. These meshes can be computed by our deformation, or be given as an input. For any given time  $t$ , our method blends both the length and the phase cross-ratios of the input



**Figure 12: Intersection-angle preserving. IAP optimization initialized with the AMAP result.**

meshes, followed by mesh reconstruction with these prescribed quantities. This is done in order to maintain a controlled amount of conformal distortion for the interpolated meshes. Our algorithm is similar in two and three dimensions, with some key differences.

**Extracting the transformation** We begin by extracting the corner reciprocals  $X$  that transform  $\mathcal{M}_0$  into  $\mathcal{M}_1$  (if not provided by our deformation). To resolve the inherent ambiguity (sign in 2D and one-parameter family per face in 3D), we solve a basic AMAP system in which the  $w$  variables are explicitly fixed to the (known) vertex positions of  $\mathcal{M}_1$ . AMAP strives to smoothly propagate the corner reciprocals, and the ambiguity is naturally resolved. Note that this works only if the two polygonal meshes are related by a piecewise-Möbius transformations (including all triangle meshes). Interpolation of non-PCM related meshes is out of our scope.

### 7.1 Interpolating the Möbius error

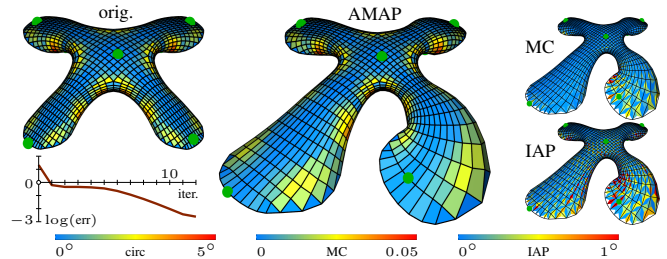
Given the corner reciprocals  $X = (cq + d)^{-1}$ , and for every inner edge  $ik$  between faces  $f, g$ , we define the Möbius error  $\Gamma_{ik}$ :

$$\Gamma_{ik} = X_{g,k} X_{f,k}^{-1} = (c_g q_k + d_g)^{-1} (c_f q_k + d_f). \quad (8)$$

And similarly for the complex  $z$  values. With compatibility conditions (4) and (7), we have that  $\Gamma_{ik} = X_{f,i} X_{g,i}^{-1}$  for complex numbers, and  $\bar{\Gamma}_{ik} = q_{ik} X_{f,i} X_{g,i}^{-1} q_{ik}^{-1}$  for quaternions. The Möbius error is multiplicative:  $\Gamma_{ik} = 1$  when two equal corner reciprocals mark a single Möbius transformation for both adjacent faces.


**Using Möbius coefficients directly** For interpolation, we do not employ vertex positions  $w$  as variables. Thus, the compatibility conditions do not suffice, since corner variables in a single polygon might not constitute an actual Möbius transformation of the face. To counter that, we employ the Möbius denominator coefficients  $c, d$  directly, and our constraints are still quadratic, as we see next.

**Relation to the cross-ratio** In two dimensions, we simply get  $(\Gamma_{ik})^2 = \text{cr}_1[i, j, k, l] / \text{cr}_0[i, j, k, l]$ . The quaternionic cross-ratio



**Figure 13: AMAP deformations of polygonal meshes.** Original (left) and AMAP (center) show concyclicity values (measured by the circle intersection angle on a diagonal of the quad) which are preserved. MC and IAP values (right) are perfect within every quad, as per a single Möbius transformation of each.





**Figure 14: Möbius error Interpolation.** The MC errors measured are bounded and are interpolated between the meshes, as evident in color and in the histograms.

does not exhibit a similar relation, but it is easy to verify that the Möbius error modulus and phase are the length and the phase cross-ratio quotients, respectively. The Möbius error is thus a Möbius invariant: it is unique up to a (right) multiplication of all  $X$  by a constant (quaternion or complex number), which is the same as a Möbius transformation. This means that mesh  $\mathcal{M}_0$  and the set of Möbius errors, define  $\mathcal{M}_1$  only up to a global Möbius transformation.

**Interpolation algorithm** Our interpolation algorithm proceeds in the following steps:

1. Extract the Möbius errors  $\Gamma_{ik}$  from meshes  $\mathcal{M}_0, \mathcal{M}_1$ .
2. Interpolate:  $\Gamma_{ik}(t) = (\Gamma_{ik})^t$ .
3. Compute interpolated Möbius coefficients  $c, d$  from the prescribed  $\Gamma_{ik}(t)$ , and reconstruct an intermediate mesh  $\mathcal{M}'(t)$ .
4. Produce the desired interpolated mesh  $\mathcal{M}(t)$  from  $\mathcal{M}'(t)$  by applying a global Möbius transformation.

There are several advantages of using this type of Möbius error interpolation: first, the length and the phase cross-ratios are interpolated gradually. Second, the Möbius error, for almost all practical cases, holds  $\text{Re}(\Gamma_{ik}) > 0$ . Furthermore, in AMAP deformations we often have  $\Gamma_{ik} \approx 1$ . Therefore, we have a meaningful logarithm of  $\Gamma_{ik}$ , and can compute the power of  $t$  unambiguously. Third, like in [Chen et al. 2013], the error is intrinsic, in the sense that it only interpolates discrete conformal quantities and not local rotations. And last, a necessary (but not sufficient) condition for a set of prescribed Möbius errors to be integrable (i.e., form integrable corner variables) is  $\prod_{k \in N(i)} \Gamma_{ik} = 1$ . This is kept under the power of  $t$ .

## 7.2 Computing the Möbius coefficients

Computing the Möbius coefficients  $c, d$  is done by prescribing the interpolated Möbius error  $\Gamma(t)$  to Equation (8), and minimizing:


	$\forall(i, k) \in \mathcal{E}, E_{\text{PRES}} =$
2D	$\sum  (c_f z_k + d_f) \Gamma_{ik}(t) - (c_g z_k + d_g) ^2 +  (c_g z_i + d_g) \Gamma_{ik}(t) - (c_f z_i + d_f) ^2$
3D	$\sum  (c_g q_k + d_g) \Gamma_{ik}(t) - (c_f q_k + d_f) ^2 +  (c_f q_i + d_f) q_{ik}^{-1} \overline{\Gamma_{ik} q_{ik}} - (c_g q_i + d_g) ^2$

subject to the quadratic reciprocal (but equivalent) versions of the edge compatibility conditions of Equations (2) and (7):

$$(c_f z_i + d_f)(z_{ik})^{-1}(c_f z_k + d_f) = (c_g z_i + d_g)(z_{ik})^{-1}(c_g z_k + d_g),$$

$$(c_f q_k + d_f)(q_{ik})^{-1}(c_f q_i + d_f) = (c_g q_k + d_g)(q_{ik})^{-1}(c_g q_i + d_g).$$

To handle the global Möbius transformation ambiguity, we first reconstruct an intermediate mesh  $\mathcal{M}'(t)$  arbitrarily, by choosing any triangle  $h$  and setting  $c_h = 0, d_h = 1$ . To initiate the optimization, we compute an unconstrained solution to the quadratic energy  $E_{\text{PRES}}$ . We get a quadratic energy in  $2(|\mathcal{F}| - 1)$  complex or quaternionic variables, and  $|\mathcal{E}|$  quadratic compatibility conditions.



**Figure 15: Interpolation of MC-equivalent meshes.** We show the quasiconformal error (MC is zero throughout), which interpolates nicely, albeit not optimized for. The arm interpolates smoothly and independently with the articulation of the fingers.

**MC error interpolation** In accordance with our framework so far,  $|\Gamma_{ik}|$  is the edge-based MC error. In order to interpolate meshes with a bounded MC error, we prescribe  $\Gamma_{ik}$  and constrain the following:

$$|c_f q_k + d_f|^2 |\Gamma_{ik}(t)|^2 = |c_g q_k + d_g|^2,$$

$$|c_g q_i + d_g|^2 |\Gamma_{ik}(t)|^2 = |c_f q_i + d_f|^2.$$

And similarly for complex  $z$ . This results in extra  $2|\mathcal{E}|$  quadratic constraints. IAP-bounded interpolation is more involved, and again not very useful in 3D, and we choose to omit it from this work.

## 7.3 Interpolating the global transformation

**Two dimensions** We choose three points on the mesh, and compute the coefficients  $a, b, c, d$  of the global initial transformation  $\mathcal{M}_0 \rightarrow \mathcal{M}_1$ . We then interpolate  $\begin{pmatrix} a(t) & b(t) \\ c(t) & d(t) \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}^t$  to obtain the transformation  $\mathcal{M}_0 \rightarrow \mathcal{M}(t)$ . In order to avoid problems with branches of matrix powers, we choose the three points in both initial meshes which are as close as possible to each other (assuming the meshes are initially aligned together), seeking a near identity. We compute the global transformation  $\mathcal{M}'(t) \rightarrow \mathcal{M}(0)$  similarly, and then compose both transformations to get  $\mathcal{M}'(t) \rightarrow \mathcal{M}(t)$ .


**Three dimensions** Since we cannot simply reposition three or four points (see the discussion in Section 5.1) we interpolate a sphere on the original mesh  $\mathcal{M}_0$  into one on  $\mathcal{M}_1$  in the following manner:

- [i] Choose three corresponding points  $p_i, p_j, p_k$  on  $\mathcal{M}_0$  and  $\mathcal{M}_1$ .
- [ii] Choose a fourth non coplanar point  $p_l$  on  $\mathcal{M}_0$  and respectively on  $\mathcal{M}_1$ , defining spheres  $S_0, S_1$ .
- [iii] Find the translations  $t \in \mathbb{H}$  and similarities  $r \in \mathbb{H}$  that transform both spheres into the origin-centered unit sphere:  $U_0 = r_0(S_0 - t_0)\bar{r}_0$ , and respectively for  $S_1$ . Transforming the corresponding three points, we get two inscribed triangles (circles) within the unit sphere.
- [iv] Find the unique Möbius transformation between the two triangles that preserves the unit sphere. It is unique by the stereographic projections of the unit sphere to the plane (such that  $p_l$  for each mesh transforms to the same halfspace), and can be found by solving the following small nonlinear problem in the Möbius coefficients  $a_u, b_u, c_u, d_u \in \mathbb{H}$ :

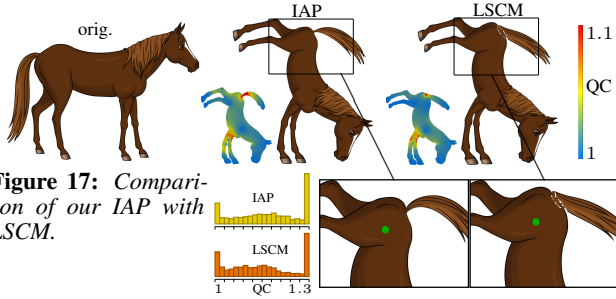
$$a_u p_{i|j|k} + b_u = c_u p_{i|j|k} + d_u, \quad |a_u p_l + b_u| = |c_u p_l + d_u|,$$

and adding imaginary-preserving conditions (6). [v] Compose the transformations to obtain  $\mathcal{M}_0 \rightarrow \mathcal{M}_1$ :

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} r_1 & t_1 \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} a_u & b_u \\ c_u & d_u \end{pmatrix} \cdot \begin{pmatrix} r_0^{-1} & -r_0^{-1} t_0 \\ 0 & 1 \end{pmatrix}.$$



**Figure 16: Interpolation of 2D MC-equivalent meshes.**



**Figure 17:** Comparison of our IAP with LSCM.

Instead of taking a matrix power, which might result in non imaginary-preserving transformations, we simply blend linearly by  $t$  with the identity matrix to achieve  $a(t), b(t), c(t), d(t)$  as before.

Mesh interpolation in three-dimensions is demonstrated in Figure 14, and MC-bounded interpolation is presented in Figures 15 and 16.

## 8 Optimizing with quadratic constraints

We use a variation of the guided projection Gauss-Newton-based algorithm of [Tang et al. 2014]. Given variables  $x$ , an energy term with the structure  $E(x) = \sum_k \alpha_k F_k^2(x)$ , and a set of constraints  $C_l(x) = 0$ , and with an initial solution  $x_0$ , we update the solution in each iteration  $i$  as follows:  $x_i = x_{i-1} + \eta_i D_i$ , where  $D_i$  is a search vector defined below. The value of  $\eta$  is determined by a bisection line search: we begin with  $\eta = 1$  and test if the maximal ( $L_\infty$ ) error of both energy and constraints is lower with  $x_i$  than with  $x_{i-1}$ . If not, we halve  $\eta$  and test again. If  $\eta < 10^{-8}$  we move to the next iteration without change. The direction  $D_i$  is found by solving  $J_i^T J_i = -J_i^T D_i$ , with the Jacobian  $J_i$ , consisting as follows:

$$J_i(x_i) = (\alpha J_F, \beta I_n, J_C)^T,$$


where  $J_F$  is the Jacobian of the terms  $F_k$ ,  $I_n$  is the identity matrix (closeness to the previous iteration), and  $J_C$  the Jacobian of the constraints  $C(x)$ . We set  $\beta = 10^{-6}$  constant in all iterations, and reduce the value of  $\alpha$  by half in each iteration, to let the system converge to the constraints. If  $\max(\alpha) < 10^{-8}$  and we are not converged, the algorithm fails, which means the results do not subscribe to the intended values (e.g., MC is not achieved). In practice, we never encounter this situation (see Figure 23). For all our energies (e.g.  $E_{AMAP}$  for deformation) we begin with  $\alpha = 10$ .

## 9 Comparisons

We compare our results with state-of-the-art relevant methods. For fair comparisons, we make small adaptations to our deformation algorithm, that we explain within context.


**2D deformations** We compare our 2D deformation to LSCM deformation [Lévy et al. 2002] in Figure 17, and BDMS [Lipman 2012] in Figure 18, as conformal deformation algorithms. Möbius transformations encourage local injectivity, avoiding drastic changes in the cross-ratio. This is evident in contrast to LSCM, where our result is locally injective, while the quasiconformal error is comparable. The comparison with BDMS is done by extracting our maximal quasiconformal error ( $QC = 1.67$ ), and using it in order to bound the distortion with the LSCM energy. Both results are locally injective with comparable error distribution.

**Conformally-equivalent triangle meshes** We show the validity of our MC transformations by reproducing our results with [Springborn et al. 2008] in the following manner (see Figure 19): we deform our mesh with positional constraints, and prescribe the resulting boundary curvature to [Springborn et al. 2008]. The result is practically indistinguishable. With a given boundary, there is a unique



**Figure 18:** Comparison of MC-optimized deformation with BDMS.

MC transformation of a domain, and so this is a good sanity check. In practice, [Springborn et al. 2008] is not defined to work with positional constraints, and the boundary curvature is prescribed directly.

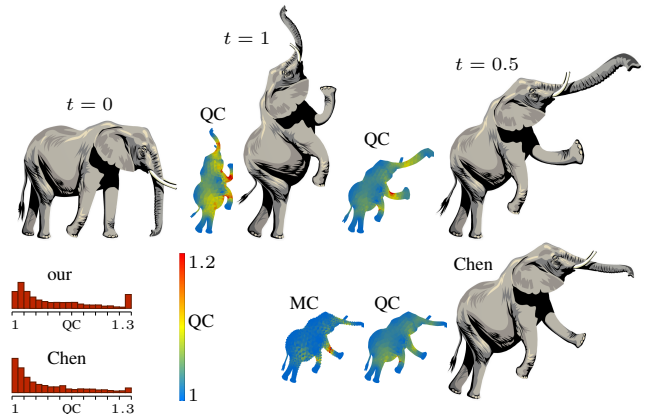


**Figure 19:** Springborn08 reproduction. Basic AMAP (center) consequently optimized for MC (right). Hausdorff distance between Springborn08 with the same boundary conditions and our MC-optimized is  $6.5 \times 10^{-6}$ .


**Planar shape interpolation with bounded distortion** Chen et al. [2013] (see Figure 20) interpolate in 2D by linearly blending squared edge lengths. In fact, our MC interpolation in 2D is equivalent to blending edge lengths logarithmically:  $l(t) = l(0)^{1-t}l(1)^t$ . This does not trivially extend to three dimensions, though. Their results blend more isometrically compared to ours, but lack the ability to precisely maintain the length cross-ratios. Both are injective and induce quite similar distribution of the quasiconformal distortion, and have bounded errors compared with the input meshes.

**Comparison with spin transformation** Since the method of spin transformations [Crane et al. 2011] does not have the ability to enforce positional constraints, we use a mesh deformed by ARAP [Sorkine and Alexa 2007] and project it onto our AMAP space. The projection was done by solving the AMAP optimization with the prescribed mesh as the initial solution, with its vertices as soft constraints (with a weighting like  $E_{AMAP}$ ). See Figure 21.

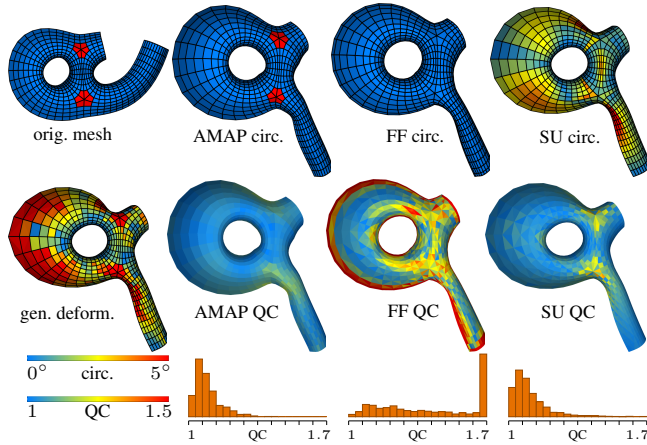
**Comparison with circular mesh deformation** In Figure 22, we compare with two methods for concyclicity-preserving deformations: Form-Finding (FF) [Tang et al. 2014] and ShapeUp [Bouaziz et al.



**Figure 20:** Comparison with bounded-distortion planar interpolation for two MC-equivalent meshes.



**Figure 21:** Comparison of AMAP with Spin Transformations. Quasiconformal error distribution is similar.



**Figure 22:** Comparison with Form-Finding (FF) and ShapeUp (SU)


2012], with the same projection mechanism as described above. Note that our PCM equivalence preserves the relative concyclicity of all faces in the mesh, while FF optimizes for it, at the price of conformality loss (ShapeUp stops at a tolerance of  $5^\circ$ ). Our result has the least quasiconformal error spread.

**Robustness** We depict the discrete nature of our formulation in Figure 23 with small and highly-constrained meshes. The Pipes circular mesh comprises 24 square faces, and 6 boundary loops, but we manage to deform and interpolate it faithfully. The Torus mesh is triangular with 50 vertices. We deform it with zero MC error.


**Quasiconformal convergence** In Figure 24 we empirically demonstrate that adhering to our framework with MC constraints can induce the convergence of the quasiconformal error in the limit. The mesh is refined by loop subdivision.

## 10 Discussion

**Implementation details** Our code and viewer environment are implemented in C++ using libigl [Jacobson and Panozzo 2014], and run on an Macbook Pro with 2.2 GHz Quad-Core Intel Core i7 processor, and 16Gb RAM. In two dimensions, for meshes with less than 4500 faces, our algorithm works interactively (less than 1 second for interaction or interpolation step). In three dimensions,



**Figure 23:** Our framework on highly-constrained configurations. Deformation and interpolation of the Pipes mesh (left), and of the torus mesh (right).



**Figure 24:** Convergence of an MC-optimized deformation with refinement (loop subdivision).

we get interactive rates with less than 3000 faces.

**Shortcomings** We empirically find our optimization to converge to the constraints, even for difficult cases. However, it is not guaranteed in theory, since our framework is nonconvex. Moreover, a rather large number of variables are employed (e.g.  $7|\mathcal{V}|$  quaternions for 3D deformation). This can be alleviated by deforming the mesh locally with most faces fixed.

**Future work** We believe our approach to discrete conformal transformations is potentially useful for many applications, such as surface flows and parametrization. We plan to study PCM transformations of tetrahedral meshes, for volumetric deformations with reduced discrete conformal distortion.

## 11 Acknowledgements

We thank Udo Hertrich-Jeromin for insightful discussions, Zohar Levi for help with comparisons, NVIDIA Corporation for the donation of the graphics card, and finally Helmut Pottmann for major support. This work was supported in part by Lise-Meitner grant M1618-N25, grant P23735-N13, and by DFG Collaborative Research Center TRR 109, ‘Discretization in Geometry and Dynamics’ through grant I 706-N26 of the Austrian Science Fund.

## References

- ALEXA, M., COHEN-OR, D., AND LEVIN, D. 2000. As-rigid-as-possible shape interpolation. In *Proc. SIGGRAPH*, 157–164.
- BEN-CHEN, M., GOTSMAN, C., AND BUNIN, G. 2008. Conformal flattening by curvature prescription and metric scaling. *Computer Graphics Forum* 27, 2, 449–458.
- BOBENKO, A. I., AND SCHRÖDER, P. 2005. Discrete Willmore flow. In *Proc. SGP*, 101–110.
- BOBENKO, A. I., AND SPRINGBORN, B. A. 2004. Variational principles for circle patterns and Koebe’s theorem. *Trans. Amer. Math. Soc* 356, 659–689.
- BOBENKO, A. I., AND SURIS, YU. S. 2008. *Discrete differential geometry. consistency as integrability*, vol. 98. Amer. Math. Soc.
- BOBENKO, A. I., HOFFMANN, T., AND SPRINGBORN, B. A. 2006. Minimal surfaces from circle patterns: geometry from combinatorics. *Ann. of Math. (2)* 164, 1, 231–264.
- BOUAZIZ, S., DEUSS, M., SCHWARTZBURG, Y., WEISE, T., AND PAULY, M. 2012. Shape-up: Shaping discrete geometry with projections. *Computer Graphics Forum* 31, 5, 1657–1667.
- CHEN, R., WEBER, O., KEREN, D., AND BEN-CHEN, M. 2013. Planar shape interpolation with bounded distortion. *ACM TOG* 32, 4 (July), 108:1–12.
- CRANE, K., PINKALL, U., AND SCHRÖDER, P. 2011. Spin transformations of discrete surfaces. *ACM TOG* 30, 4, 104:1–10.



FRENKEL, I., AND LIBINE, M. 2008. Quaternionic analysis, representation theory and physics. *Adv. Math.* 218, 6, 1806–1877.

GU, X., AND YAU, S.-T. 2002. Global conformal surface parameterization. In *Proc. SGP*, 127–137.

HABBECKE, M., AND KOBELT, L. 2012. Linear analysis of non-linear constraints for interactive geometric modeling. *Computer Graphics Forum* 31, 641–650.

HERTRICH-JEROMIN, U. 2003. *Introduction to Möbius differential geometry*, vol. 300 of *London Mathematical Society Lecture Note Series*. Cambridge University Press.

JACOBSON, A., AND PANOZZO, D., 2014. libigl: A simple C++ geometry processing library. <http://libigl.github.io/libigl/>.

JAKOBS, W., AND KRIEG, A. 2010. Möbius transformations on  $\mathbb{R}^3$ . *Complex Variables and Elliptic Equations* 55, 4, 375–383.

KHAREVYCH, L., SPRINGBORN, B. A., AND SCHRÖDER, P. 2006. Discrete conformal mappings via circle patterns. *ACM TOG* 25, 2, 412–438.

KILIAN, M., MITRA, N. J., AND POTTMANN, H. 2007. Geometric modeling in shape space. *ACM Trans. Graph.* 26, 3, 1–8.

LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least squares conformal maps for automatic texture atlas generation. *ACM TOG* 21, 3, 362–371.

LIPMAN, Y., COHEN-OR, D., GAL, R., AND LEVIN, D. 2007. Volume and shape preservation via moving frame manipulation. *ACM TOG* 26, 1, 5:1–14.

LIPMAN, Y. 2012. Bounded distortion mapping spaces for triangular meshes. *ACM TOG* 31, 4 (July), 108:1–13.

LIU, Y., POTTMANN, H., WALLNER, J., YANG, Y.-L., AND WANG, W. 2006. Geometric modeling with conical meshes and developable surfaces. *ACM TOG* 25, 3, 681–689.

LIU, L., ZHANG, L., XU, Y., GOTSMAN, C., AND GORTLER, S. J. 2008. A local/global approach to mesh parameterization. In *Proc. SGP*, 1495–1504.

MÜLLER, C. 2011. Conformal hexagonal meshes. *Geometriae Dedicata* 154, 27–41.

PARIES, N., DEGENER, P., AND KLEIN, R. 2007. Simple and efficient mesh editing with consistent local frames. In *Proc. PG*, 461–464.

POTTMANN, H., LIU, Y., WALLNER, J., BOBENKO, A. I., AND WANG, W. 2007. Geometry of multi-layer freeform structures for architecture. *ACM TOG* 26, 3, 65:1–11.

SCHÜLLER, C., KAVAN, L., PANOZZO, D., AND SORKINE-HORNUNG, O. 2013. Locally injective mappings. *Computer Graphics Forum* 32, 5, 125–135.

SORKINE, O., AND ALEXA, M. 2007. As-rigid-as-possible surface modeling. In *Proc. SGP*, 109–116.

SPRINGBORN, B. A., SCHRÖDER, P., AND PINKALL, U. 2008. Conformal equivalence of triangle meshes. *ACM TOG* 27, 3, 77:1–11.

TANG, C., SUN, X., GOMES, A., WALLNER, J., AND POTTMANN, H. 2014. Form-finding with polyhedral meshes made simple. *ACM TOG* 33, 4, 70:1–9.

VAXMAN, A. 2014. A projective framework for polyhedral mesh modelling. *Computer Graphics Forum* 33, 8, 121–131.

WEBER, O., AND ZORIN, D. 2014. Locally injective parametrization with arbitrary fixed boundaries. *ACM TOG* 33, 4, 75.

WEBER, O., MYLES, A., AND ZORIN, D. 2012. Computing extremal quasiconformal maps. In *Computer Graphics Forum*, vol. 31, Wiley Online Library, 1679–1689.

WILKER, J. 1993. The quaternion formalism for Möbius groups in four or fewer dimensions. *Linear Algebra Appl.* 190, 99–136.

WINKLER, T., DRIESEBERG, J., ALEXA, M., AND HORMANN, K. 2010. Multi-scale geometry interpolation. *Computer Graphics Forum* 29, 2, 309–318.

YANG, Y.-L., YANG, Y.-J., POTTMANN, H., AND MITRA, N. J. 2011. Shape space exploration of constrained meshes. *ACM TOG* 30, 6, #124,1–12.

## A Imaginary-preserving transformations

**Imaginary-preserving conditions (Equation (6))** We look for sets  $a, b, c, d \in \mathbb{H}$  s.t.  $\forall q = (0, v_q) \in \text{Im } \mathbb{H}$  we have  $(aq + b)(cq + d)^{-1} \in \text{Im } \mathbb{H}$ . Developing:

$$(aq + b)(cq + d)^{-1} = (aq + b)(\bar{q}\bar{c} + \bar{d})/|cq + d|^2.$$

The norm in the denominator is scalar and thus irrelevant. We obtain:

$$(aq + b)(\bar{q}\bar{c} + \bar{d}) = |q|^2 a\bar{c} + b\bar{d} + aq\bar{d} + b\bar{q}\bar{c} \in \text{Im } \mathbb{H}.$$

For this expression to be imaginary for all  $q$ , we see that  $a\bar{c}, b\bar{d} \in \text{Im } \mathbb{H}$  is required. In vector notation, the other terms spell:

$$0 = \text{Re}(aq\bar{d} + b\bar{q}\bar{c}) = \text{Re}(\bar{d}aq + \bar{c}bq) = \langle v_{\bar{d}a}, v_q \rangle + \langle v_{\bar{c}b}, v_q \rangle,$$

for all  $v_q \in \mathbb{R}^3$ . Consequently,  $v_{\bar{d}a} + v_{\bar{c}b} = 0$  and thus  $\bar{a}d - \bar{b}c \in \mathbb{R}$ .

**Imaginary-preserving edge transformation** For general quaternionic Möbius transformation  $m(q) = (aq + b)(cq + d)^{-1}$  s.t.  $\begin{pmatrix} a' & b' \\ c' & d' \end{pmatrix} := \begin{pmatrix} a & b \\ c & d \end{pmatrix}^{-1}$  we get [Frenkel and Libine 2008]:

$$\begin{aligned} m(q_k) - m(q_i) &= (a' - q_i c')^{-1} (q_k - q_i) (c q_k + d)^{-1} \\ &= (a' - q_k c')^{-1} (q_k - q_i) (c q_i + d)^{-1}. \end{aligned}$$

Imaginary conditions (6) give a simple expression for the inverse

$$\begin{pmatrix} a' & b' \\ c' & d' \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}^{-1} = (\bar{a}d + \bar{c}b)^{-1} \begin{pmatrix} \bar{d} & \bar{b} \\ \bar{c} & \bar{a} \end{pmatrix},$$

as easily verified by multiplying from the right by  $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ . Note that  $\bar{a}d - \bar{b}c \in \mathbb{R} \Leftrightarrow \bar{a}d + \bar{c}b \in \mathbb{R}$ . Consequently, multiplying from the left hand side and the uniqueness of the inverse implies  $\bar{a}d + \bar{c}b = \bar{a}d + \bar{b}c$ . We mention that the conditions  $\bar{a}d + \bar{b}c \in \mathbb{R}$ ,  $\bar{a}\bar{b}, \bar{c}\bar{d} \in \text{Im } \mathbb{H}$  are equivalent to (6).