

Canonical Möbius Subdivision

AMIR VAXMAN, Utrecht University, Netherlands

CHRISTIAN MÜLLER, TU Wien, Austria

OFIR WEBER, Bar-Ilan University, Israel




Fig. 1. An efficient coarse-to-fine architectural design using our method. A coarse control cage (*left*) is edited by the green handles, and optimized for Möbius regularity (*top-center*; front and rear perspectives) using the optimization of [Vaxman et al. 2017], producing a mesh that is as spherical as possible (taking 1.0 seconds). The mesh is subdivided (0.5 seconds) into a fine polygonal pattern, retaining the spherical features (*bottom-center*; front and rear perspectives). *Right*: A zoom-in. A similar result by direct optimization on the fine mesh would take more than a minute.

We present a novel framework for creating Möbius-invariant subdivision operators with a simple conversion of existing linear subdivision operators. By doing so, we create a wide variety of subdivision surfaces that have properties derived from Möbius geometry; namely, reproducing spheres, circular arcs, and Möbius regularity. Our method is based on establishing a canonical form for each 1-ring in the mesh, representing the class of all 1-rings that are Möbius equivalent to that 1-ring. We perform a chosen linear subdivision operation on these canonical forms, and blend the positions contributed from adjacent 1-rings, using two novel Möbius-invariant operators, into new face and edge points. The generality of the method allows for easy coarse-to-fine mesh editing with diverse polygonal patterns, and with exact reproduction of circular and spherical features. Our operators are in closed-form and their computation is as local as the computation of the linear operators they correspond to, allowing for efficient subdivision mesh editing and optimization.

CCS Concepts: • **Computing methodologies** → **Mesh models**; **Mesh geometry models**;

Additional Key Words and Phrases: Möbius transformations, mesh subdivision, conformal transformations, regular meshes, architectural geometry

Authors' addresses: Amir Vaxman, Utrecht University, Princetonplein 5, De Uithof, Utrecht, 3584CC, Netherlands, a.vaxman@uu.nl; Christian Müller, TU Wien, Wiedner Hauptstr. 8-10/104, Vienna, A-1040, Austria, cmueller@geometrie.tuwien.ac.at; Ofir Weber, Bar-Ilan University, Building 1102, Room 404, Ramat Gan, 52900, Israel, ofir.weber@biu.ac.il.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

0730-0301/2018/11-ART227 \$15.00

<https://doi.org/10.1145/3272127.3275007>

ACM Reference Format:

Amir Vaxman, Christian Müller, and Ofir Weber. 2018. Canonical Möbius Subdivision. *ACM Trans. Graph.* 37, 6, Article 227 (November 2018), 15 pages. <https://doi.org/10.1145/3272127.3275007>

1 INTRODUCTION

Two crucial measures of mesh design are the quality of the resulting shape and the efficiency of the design process. The quality of the shape can be measured by its overall smoothness, and by the properties of individual mesh elements, such as their symmetry and regularity. The design process should support the editing of the shape in several levels of detail. It should adhere to user-specified constraints, and should be interactive. Finally, the designer would like to work with a diverse set of possible shapes and mesh patterns, to get a creative edge to her work.

A popular design paradigm is coarse-to-fine editing with *subdivision surfaces*. This comprises working with a coarse mesh (serving as a *control mesh*), and obtaining smoother meshes by the process of recursive subdivision of mesh elements. The nested representation makes it also possible to encode high-frequency details in each level. The most commonly used methods for creating the subdivision hierarchy are linear, as they are optimally efficient to compute and with theoretical guarantees of smoothness.

Many applications require extra properties from the result. For instance, adherence to user constraints or reproduction of geometries like spheres or planes. The straightforward solution for this is “subdivide-and-optimize”, where the method interleaves linear subdivision with nonlinear costly optimization.

An alternative is creating nonlinear methods that reproduce desired geometric features by construction. We offer a general method to modify linear subdivision methods locally, such that the subdivision would be invariant to Möbius transformations. As such, the




Fig. 2. The designed and subdivided mesh from Fig. 1 presented in an architectural context, as a cover for a metro station.

subdivision reproduces spheres and circular arcs and, in general, produces surfaces with relatively low Willmore energy with unique aesthetic features. The modification is done by transforming each 1-ring to a canonical form on which we apply a linear subdivision operator. The points shared by several 1-rings are then transformed back to the original mesh and blended with Möbius-invariant operators. With this, we are able to convert any linear subdivision scheme into a Möbius-invariant one, including operators that create unconventional polygonal patterns. Since our conversion is closed-form and local, its asymptotic computational complexity is linear in the number of mesh elements, just like linear subdivision methods, making it suitable for interactive multiresolution editing and design.

2 RELATED WORK

2.1 Subdivision surfaces

Linear and stationary methods. The “classic” subdivision schemes in both computer-aided design and geometry processing are both linear (described as a matrix multiplication operator), and stationary (the subdivision stencil is constant). Popular schemes are Catmull-Clark [Catmull and Clark 1978] and Kobbelt [Kobbelt 1996] for quadrilateral meshes, and Loop [Loop 1987] and Butterfly [Dyn et al. 1990] for triangle meshes. In addition, Akleman et al. [2005] presented methods for creating more general polygonal patterns from simple subdivision operators. It is beyond the scope of this paper to discuss the plethora of such methods and their properties, and we refer the interested reader to dedicated texts such as [Cashman 2012] and [Cavaretta et al. 1991].

The advantage of linear subdivision schemes is that in addition to being relatively smooth, they are simple and efficient to evaluate and have good convergence properties. As such, we choose them as the basic ingredients from which we build our Möbius-invariant subdivision operators.

Nonlinear subdivisions. In recent years, a few works presented nonlinear schemes as extensions to (or conversions of) the linear

approaches, often for non-trivial geometries. By doing so, they could inherit many of the smoothness and convergence properties in a provable way. Notable examples are the log-exp scheme [Rahman et al. 2005] for manifold-valued data, and geodesic averages [Wallner and Dyn 2005]. Some schemes are build to reproduce specific geometries and functions. For example, exponentials of linear functions are reproduced in [Micchelli 1996], and convex functions are reproduced with the harmonic mean for interpolatory schemes in [Floater and Micchelli 1998].

Two nonlinear methods are of particular relation to our method. The first is the method of [Sabin and Dodgson 2005] which reproduces circles with the four point scheme construction. Our $F4$ operator (Section 4.4.1) is conceptually similar, replacing the four-point scheme in a way that allows us to generalize the construction to surfaces lying on spheres, and it is Möbius-invariant. The second related method is that of [Schaefer et al. 2008]. They show that one can reproduce diverse geometric properties by conjugating a linear operator with a nonlinear transformation. For instance, a circle can be reproduced by locally transforming points to a line by a Möbius transformation. As such, they share the idea of subdivision in a canonical form (albeit only for curves). However, the Möbius transformation needs to be specified to create the rule, and only the specific choice that brings a circle to a line would reproduce a circle by refinement. In contrast, we devise the canonical form by Möbius invariants such as the cross-ratio, and therefore always reproduce spheres and circles by construction.

2.2 Discrete Möbius geometry processing

Architectural geometry. Objects of Möbius geometry, namely generalized spheres and circles, gained popularity in geometric design propelled by applications in architectural geometry [Pottmann et al. 2015]. For instance, circular meshes, where every face is circumscribed by a circle, have vertex offsets [Pottmann et al. 2007], allowing them to be realized with nodes without torsion. Circular arc structures [Bo et al. 2011] use parts of circles in lieu of straight edges, and are also Möbius-invariant. However, we do not know of any algorithm that targets exact reproduction of spheres in architectural design, such as ours. We note that there are algorithms that are capable of handling generic geometric constraints, such as sphere reproduction, by projection. As such, one can apply the “subdivide-and-optimize” metaphor for level-of-detail sphere and circle reproduction, by interleaving subdivision with nonlinear projection. Some methods provide a general framework for such projections; for instance, Shape-up [Bouaziz et al. 2012] and quadratic constraint projection [Tang et al. 2014]. However, they rely on nonlinear global optimization, which doesn’t scale well to finely-subdivided meshes. Our method specifically targets Möbius-invariant properties of meshes, and therefore is not as generic. However, it is local and in closed-form and therefore far more efficient to compute.

Willmore and regular meshes. The smooth Willmore energy of a surface is given by $\int_S (\kappa_1 - \kappa_2)^2 dA$, where κ_1, κ_2 are the principal curvatures. It is a measure of how much a given surface S deviates from a sphere, and it is known to be invariant to Möbius transformations. Willmore meshes have been studied in geometry processing recently. They are obtained as critical points of a discrete analogue

of the Willmore energy. Crane et al. [2013] define a discrete Willmore flow, within the framework of conformal spin transformations, for the purpose of surface fairing. In [Bobenko and Schröder 2005], a discrete Willmore energy is devised by looking at the intersection angles of circumscribed circles of the triangles of a mesh. This definition preserves the Möbius invariance of the continuous Willmore energy. Our method uses the latter definition, and modifies subdivision operators to favor meshes with low Willmore energy, due to its Möbius-invariant nature.

Vaxman et al. [2017] introduce the definition of *Möbius-regularity*. A mesh is Möbius-regular if every 1-ring is equivalent to a symmetric 1-ring whose faces are regular polygons by a Möbius transformation. This definition is closely related to that of Willmore meshes, since both definitions advocate embedding of local neighborhoods in spheres. Nevertheless, as demonstrated in [Vaxman et al. 2017], Möbius-regularity enforces additional structure. Our subdivision operator is designed to preserve Möbius-regularity, i.e., given a Möbius-regular mesh as input, it will produce a finer mesh which is also Möbius-regular.

Polygonal patterns. There is a recent interest in polygonal patterns [Akleman et al. 2005; Jiang et al. 2017, 2015; Peng et al. 2018]. Like circular meshes, it is motivated by applications in design, trying to mimic beautiful geometric patterns found in Arabesques and modern art, or for the purpose of efficient topological optimizations. Given the linear subdivision operators by [Akleman et al. 2005], our algorithm naturally works with spherical and near-spherical pattern design; see Figures 1 and 2 for an example.

3 MÖBIUS GEOMETRY PROCESSING

The subdivision setting that we present uses notions related to discrete conformal geometry and Möbius transformations as recently used in [Vaxman et al. 2017]. This work is based on definitions and methods from *discrete differential geometry* such as circle patterns [Bobenko and Springborn 2004; Kharevych et al. 2006], discrete conformality [Springborn et al. 2008] and the discrete Willmore energy [Bobenko and Schröder 2005]. For brevity, we only provide the necessary details needed for this paper, and refer the reader to the aforementioned sources for a more in-depth treatment.

3.1 Preliminaries and notation

We apply our subdivision to 2-manifold meshes $\mathcal{M} = (\mathcal{V}, \mathcal{E}, \mathcal{F})$, with or without boundaries, in \mathbb{R}^3 . We support pure triangular meshes (where all faces are triangles), pure quad meshes, and general polygonal meshes. A *vertex star* vu_1, \dots, vu_n consists of all vertices u_1, \dots, u_n that are connected along an edge to a *central vertex* v . The vertices u_1, \dots, u_n are assumed to be in cyclic order. We call the n -gon that connects the vertices u_i the *boundary polygon* and denote it by B_v . Note that if the mesh is not triangular, the boundary polygon B_v is different from the boundary of the 1-ring as the 1-ring consists of all the faces around the vertex v . We illustrate this in Figure 3.

Quaternions for geometry in \mathbb{R}^3 . It is convenient to represent objects and transformations in three-dimensional Möbius geometry




Fig. 3. A vertex star vu_i with boundary polygon B_v . For triangle meshes (left), the boundary polygon $B_v = u_1, \dots, u_6$ consists of actual edges of the mesh. For polygonal meshes (right), the boundary polygon $B_v = u_1, u_2, u_3$, is the boundary of the vertex star (red dashed lines). The polygon r_1, \dots, r_9 is the boundary of the 1-ring (green lines). The boundary of the 1-ring and the boundary polygon are different in this case.

using quaternions. *Quaternions* $q \in \mathbb{H}$ are “four-dimensional numbers” $q = [r, v]$ with a one-dimensional *real part* $r = \text{Re}(q) \in \mathbb{R}$ and a three-dimensional *imaginary part* $v = \text{Im}(q) \in \text{Im } \mathbb{H} \cong \mathbb{R}^3$. We represent vertex positions (x, y, z) in \mathbb{R}^3 by imaginary quaternions $[0, (x, y, z)]$ and the edge vector between two points $q_i, q_j \in \text{Im } \mathbb{H}$ is denoted by $q_{ij} = q_j - q_i$. Quaternions can be added $[r, v] + [s, w] = [r+s, v+w]$ and multiplied $[r, v] \cdot [s, w] = [rs - \langle v, w \rangle, rw + sv + v \times w]$ ($\langle \cdot, \cdot \rangle$ and \times stand for Euclidean inner product and cross product, respectively). The multiplication is commutative if and only if (v, w) are linearly dependent. We denote the conjugate by a bar: $\overline{[r, v]} = [r, -v]$ with the property $\overline{\overline{q}} = q$, which immediately implies $\overline{q} = -q$ for all $q \in \text{Im } \mathbb{H}$. A quaternion q has an absolute value $|q| = \sqrt{q\overline{q}} \in \mathbb{R}$. It has an inverse $q^{-1} = \overline{q}/|q|^2$ if $q \neq 0$. Any non-real quaternion q can be represented in its polar form $q = |q|[\cos \theta, n \sin \theta]$ with $\|n\| = 1$. To dispel ambiguity in the angle, we always use the convention $\theta \in (-\pi, \pi]$. As such, there is a well-defined square root, given by $\sqrt{q} = \sqrt{|q|}[\cos \frac{\theta}{2}, n \sin \frac{\theta}{2}]$. Note that the square root always has a non-negative real part by this choice. Finally, for any quaternion $q \neq 0$, the transformation $x \mapsto q^{-1}xq$ produces a quaternion whose imaginary part is a Euclidean rotation of the imaginary part of x , and whose real part equals that of x , i.e. $|\text{Im}(x)| = |\text{Im}(q^{-1}xq)|$ and $\text{Re}(x) = \text{Re}(q^{-1}xq)$.

3.2 Möbius transformations and the cross-ratio

Möbius transformations in quaternions. We use the term *generalized sphere* to denote an ordinary sphere or a plane (which can be interpreted as a sphere centered at ∞). Möbius transformations in \mathbb{R}^3 are bijective maps from $\text{Im } \mathbb{H} \cup \{\infty\}$ to itself. They can be represented concisely using quaternions:

$$q \mapsto w = (aq + b)(cq + d)^{-1}, \quad (1)$$

with a suitable choice of coefficients $a, b, c, d \in \mathbb{H}$ that guarantees that $w \in \mathbb{R}^3$ [Vaxman et al. 2015]. Any Möbius transformation can be interpreted as a composition of rigid transformations, scaling, and inversions in *generalized spheres*. As such, Möbius transformations transform generalized spheres into generalized spheres, and generalized circles (circles and lines) into generalized circles. Two sets of points are called *Möbius equivalent* if there is a Möbius transformation that maps one set to the other.


Cross-ratio. The *quaternionic cross-ratio* of four points, $q_i, q_j, q_k, q_l \in \text{Im } \mathbb{H}$ is defined by:

$$\text{cr}[q_i, q_j, q_k, q_l] = q_{ij} q_{jk}^{-1} q_{kl} q_{li}^{-1}.$$

The quaternionic cross-ratio is real if and only if the four points are cocircular [Bobenko and Pinkall 1996]. Unlike the complex cross-ratio, which is fully invariant under Möbius transformations, the quaternionic cross-ratio is not fully preserved. When we transform four points using the Möbius transformation in Eq. (1), the quaternionic cross-ratio cr is transformed by a rotation:

$$\text{cr}[w_i, w_j, w_k, w_l] = (cq_i + d) \text{cr}[q_i, q_j, q_k, q_l] (cq_i + d)^{-1}.$$

Hence, by the rotation property mentioned at the end of Section 3.1, we have that $\text{Re}(\text{cr})$, $|\text{Im}(\text{cr})|$ are preserved by Möbius transformations but the direction of $\text{Im}(\text{cr})$ is not preserved in general. The imaginary part of the cross-ratio has a meaningful geometric interpretation: it has the same direction as the 3D vector that connects the center of the circumsphere of q_i, q_j, q_k, q_l to the vertex q_i (see inset). As such, the two cross-ratios $\text{cr}[q_i, q_j, q_k, q_l]$ and $\text{cr}[q_k, q_l, q_i, q_j]$ can be parallel only when the radius of the circumsphere is infinite (and then it is a plane), or when q_i and q_k are antipodal points. We next use this cross-ratio characterization of spheres and circles to define our discrete Willmore energy and Möbius regularity.



3.3 Tangent polygons, Willmore energy, and Möbius regularity

A fundamental element in our construction is the *tangent polygon* of a vertex star. The tangent polygon is explored in [Vaxman et al. 2017], where it is used to characterize cosphericality and regularity of mesh elements. We repeat important concepts of the characterization given by [Vaxman et al. 2017] for completeness, and extend them with several novel insights.

Corner tangents and tangent polygons. The *corner tangent* of three points $q_i, q_j, q_k \in \text{Im } \mathbb{H}$ is defined as:

$$t[q_i, q_j, q_k] := q_{ij}^{-1} + q_{jk}^{-1}.$$

Geometrically, the corner tangent is a 3D vector that is tangent to the circumcircle of the three points q_i, q_j, q_k at the corner point q_j (as illustrated in Fig. 4 left). Given a vertex star vu_1, \dots, vu_n , we can sum the corner tangents around it. The closedness of the star then implies that the sum must vanish:

$$\sum_i t[u_i, v, u_{i+1}] = \sum_i (v - u_i)^{-1} + (u_{i+1} - v)^{-1} = 0.$$

Hence, we can interpret the corner tangent vectors associated with vertex v as the edges of a *closed* abstract n -gon, which is denoted as T_v . This is the so-called *tangent polygon*. In order to define T_v uniquely avoiding translation ambiguity, we always refer in this paper to one particular instance of T_v whose vertices are written in the form $T_i = (u_i - v)^{-1}$, in which case we get that $t[u_i, v, u_{i+1}] = T_{i+1} - T_i$ (see Fig. 4 right). The following Lemma is one of our key contributions, as it leads to several further insights. It can be used

to justify most of the properties of T_v that are explored in [Vaxman et al. 2017] in an intuitive and concise manner.

LEMMA 3.1. *Let vu_1, \dots, vu_n be a vertex star. Then B_v and T_v are Möbius equivalent.*

PROOF. The Möbius transformation $q \mapsto (q - v)^{-1}$ maps every u_i to $(u_i - v)^{-1} = T_i$, and with that the boundary polygon B_v is Möbius-equivalent to the tangent polygon T_v . \square




Fig. 4. A vertex star (left) with one corner tangent at v , and the complete resulting closed tangent polygon T_v (right). Note that the tangent polygon is not necessarily planar. The imaginary part of the cross-ratio $\text{cr}[v, u_3, u_4, u_5]$ is orthogonal to the two corresponding edges of the tangent polygon (right).

Given four points q_i, q_j, q_k, q_l as before, it is possible [Vaxman et al. 2017, Section 4.3] to express their cross-ratio $\text{cr}[q_i, q_j, q_k, q_l]$ in terms of corner tangents as follows:

$$\text{cr}[q_i, q_j, q_k, q_l] = t[q_k, q_i, q_j]^{-1} \cdot t[q_k, q_i, q_l]. \quad (2)$$

This immediately implies that the imaginary part of the cross-ratios $\text{cr}[v, u_{i-1}, u_i, u_{i+1}]$ is orthogonal to the edges $T_{i-1}T_i$ and T_iT_{i+1} of the tangent polygon (see Fig. 4).

The Möbius transformation $q \mapsto w = (aq + b)(cq + d)^{-1}$ maps the corner tangent $t[u_i, v, u_{i+1}]$ to:

$$t[w_i, w, w_{i+1}] = (cv + d) t[u_i, v, u_{i+1}] (\overline{cv + d}),$$

which corresponds to a similarity. We thus obtain the following result:

LEMMA 3.2. *The tangent polygons of Möbius equivalent vertex stars are similar to each other (see Fig. 5).*

Discrete Willmore energy. Consider a vertex star vu_i as before, and its tangent polygon T_v . Recall that the imaginary parts of the cross-ratios $\text{cr}[v, u_{i-1}, u_i, u_{i+1}]$ are codirectional with the vector between the center of the circumsphere to these points and v . As such, if and only if all $\text{Im}(\text{cr}[v, u_{i-1}, u_i, u_{i+1}])$ are codirectional, which means that T_v is planar, the entire vertex star is cospherical.

Hence, one can measure local deviation of a 1-ring from being cospherical by measuring the planarity of the tangent polygon instead, which is often easier. This serves as a discrete *Willmore energy* for the vertex star. [Bobenko and Schröder 2005] used an equivalent measure of the angle defect of the tangent polygon for that purpose. We note that this notion of discrete Willmore energy is rather weak compared to the continuous one. Consider for instance

a mesh with random vertex positions, and then project each vertex onto a unit sphere by normalization: $v_i \rightarrow v_i/|v_i|$. Such a mesh will have zero discrete Willmore energy despite its lack of regularity and smoothness; all the tangent polygons would be planar, but severely distorted and self intersecting.

Möbius regularity and regular canonical forms. In the case where the tangent polygon of a vertex star is *regular*, i.e., a perfectly Euclidean regular n -gon, we call the original vertex star *Möbius regular* (see Fig. 5). This means that this vertex star is Möbius equivalent to a Euclidean regular vertex star, i.e., a central vertex connected to the vertices of a regular n -gon. This perfect star is called the *regular canonical form*. *Möbius regular faces* on the other hand are Möbius equivalent to Euclidean regular n -gons. A *Möbius regular mesh* is then consequently a mesh where all faces and all vertex stars are Möbius regular. We note that in the case where a 1-ring is *mixed*, adjacent to polygonal faces of different valences, a Möbius-regular vertex star is Möbius equivalent to a canonical form which is not perfectly regular, but still *cocircular*, and where the edge lengths depend on the valences of the respective faces.

Our paper generalizes the definition of canonical forms to *all* 1-rings, not just Möbius-regular 1-rings with regular canonical forms. We rely on the relation between vertex stars and their tangent polygons for our subdivision method (Section 4). By working on the canonical forms instead of on the original vertex stars, and using linear schemes that preserve planarity and Euclidean regularity, we preserve the cosphericality and the Möbius-regularity of the original meshes in the respective elements.




Fig. 5. Left: a Euclidean-regular vertex star (green) with its tangent polygon (red). Right: a Möbius-regular vertex star (green) with its Euclidean regular tangent polygon (red). The two tangent polygons of Möbius equivalent vertex stars are related by a similarity.

4 CANONICAL MÖBIUS SUBDIVISION

In the following, we describe a general method to convert linear subdivision methods into new subdivision methods that commute with Möbius transformations. That is, subdividing two coarse Möbius equivalent meshes leads to two fine meshes which are related by the same Möbius transformation. An effect of this is that linear methods that preserve planarity (basically all of them) are converted into methods that preserve cosphericality and methods that preserve Euclidean regularity in the plane (all the methods we use) into methods that preserve Möbius regularity.

4.1 The subdivision algorithm

Our subdivision algorithm proceeds in the following steps:

- For each vertex v , compute a canonical form C_v and a Möbius transformation m_v from the original star to the canonical form. Transform the entire 1-ring by this transformation (Sec. 4.2).
- Create new *candidate* edge, face, and vertex points with chosen linear stencils in the canonical form C_v (Sec. 4.3).
- Transform the new candidate points back to the original mesh via the inverse transformation m_v^{-1} .
- Blend candidate points that are computed from all adjacent vertex 1-rings with Möbius invariant blending operators into final edge and face points (Sec. 4.4).

Because of the use of a canonical form, we denote this as *canonical Möbius subdivision*. The generality of this scheme allows us to use *any* linear scheme that can be described by blending points from individual vertex stars. The process is illustrated in Figure 6.

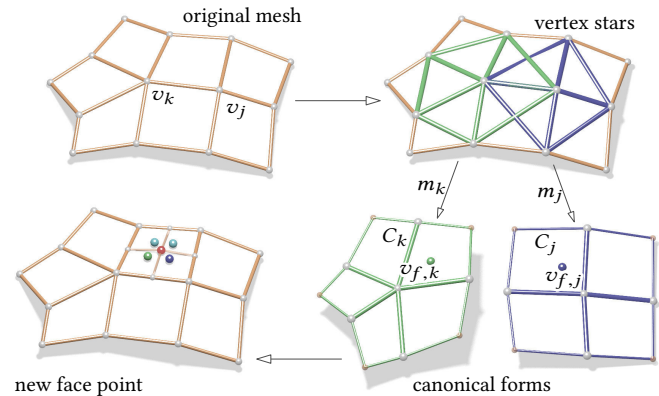


Fig. 6. Illustrating the canonical subdivision process: we transform the vertex stars of an original mesh into canonical forms C_j and C_k , where we compute new face candidate points in face f , namely $v_{f,j}$ and $v_{f,k}$, with a linear subdivision scheme. Then, we transform the points back, and blend them across stars into the final face point v_f (red).

4.2 Canonical forms

The first step of our algorithm is to transform each 1-ring to a canonical form, on which the linear subdivision operators would perform. As such, a canonical form should be a 1-ring of the same connectivity as the original, and it should be invariant up to similarities when applying a Möbius transformation to the original. Furthermore, we need to construct a Möbius transformation (and consequently its inverse) between the two 1-rings.

A similar construction to what we call *canonical forms* has been introduced as *canonical embedding* for Möbius-regular meshes in [Vaxman et al. 2017] (see Section 3.3). The canonical embeddings serve as representatives of the equivalence class of Möbius-regular vertex stars. However, our definition extends this to a canonical form for equivalence classes of *every* type of vertex stars, not just Möbius-regular ones.

Consider a vertex star vu_1, \dots, vu_n around central vertex v with boundary polygon $B_v = u_1, \dots, u_n$ as before, for which we compute the canonical form C_v (see Figure 7). The natural candidate for the

boundary of the canonical form is based on the tangent polygon T_v , which is Möbius equivalent to B_v (Lemma 3.1), and which is the same up to similarities for all Möbius-equivalent vertex stars (Lemma 3.2).




Fig. 7. Constructing the canonical form. *Left*: a given vertex star with central vertex v and boundary polygon $B_v = u_1, \dots, u_n$. *Center*: the tangent polygon T_v with vertices $T_i = m_{B_v \rightarrow T_v}(u_i)$ and its approximate best-fit circle with barycenter $c(T_v)$. The central vertex v gets mapped to infinity: $m_{B_v \rightarrow T_v}(v) = \infty$. *Right*: inversion in the best-fit circle gives the canonical form $C_v = m_{T_v \rightarrow C_v}(T_v)$ with $c(T_v) = m_{T_v \rightarrow C_v}(\infty)$.

However, the Möbius transformation $m_{B_v \rightarrow T_v}(q) = (q - v)^{-1}$ sends v to ∞ . That means that it is not sufficient for constructing a canonical form. To complete our construction, we compose $m_{B_v \rightarrow T_v}$ with another transformation $m_{T_v \rightarrow C_v}(q)$ that brings ∞ into a point $c(T_v)$ that we consider as the center of T_v . There are two nested requirements on $c(T_v)$: if T_v is planar, then $c(T_v)$ needs to be in the plane containing T_v . This will guarantee spherical reproduction. If T_v is in addition a regular n -gon, $c(T_v)$ needs to be exactly the center of its circumscribing circle. This will guarantee Möbius-regularity reproduction. To adhere to these requirements, it is possible to compute a best-fit circle to the projection of T_v to a 2D plane, and use the center of the circle as $c(T_v)$. We denote the radius of the circle as $r(T_v)$; then, to complete the computation of the canonical form, we define $m_{T_v \rightarrow C_v}(q) = r^2(c - q)^{-1} + c$. In words, $m_{T_v \rightarrow C_v}$ is an inversion of T_v in (the sphere whose equator is) its own best-fit circle, which brings ∞ to $c(T_v)$, and preserves T_v if it is indeed regular (or just cocircular).

Nevertheless, in practice we avoid a costly nonlinear (albeit local) computation of best-fit circle, and instead approximate $c(T_v)$ by simply using the barycenter of T_v , and $r(T_v)$ as the average distance of $c(T_v)$ from the points of T_v . The barycenter coincides with the exact $c(T_v)$ for regular n -gons, so we do not harm the desired Möbius regularity reproduction.

Finally, the complete transformation from the vertex star to its canonical form is the Möbius transformation:

$$m_{B_v \rightarrow C_v} := m_{T_v \rightarrow C_v} \circ m_{B_v \rightarrow T_v}.$$

Boundary canonical forms. In the general case, boundary vertices also have a tangent polygon, which is an open polygon. For Möbius-regular boundary vertices, it is part of a regular polygon. To generalize the construction above, we need to compute $c(T_v)$ differently. In practice, we found that a good choice is simply the average of the two neighboring boundary vertices. Note it still reproduces regularity when the boundary has a regular valence (i.e., 3 edges for quad meshes and 4 edges for triangle meshes).

A special case is boundary “ear” vertices of valence 2. Ear vertices are adjacent to a single face. As such, we use all the vertices of the




Fig. 8. Polygonal pattern subdivision with our method. Note how the tip of the cone subdivides into a near-spherical pattern, even with unconventional pattern subdivision.

face for the construction of B_v of the ear vertex. This reproduces our desired properties as well.

4.3 Linear subdivision schemes

We work with a broad range of linear subdivision schemes designed for triangle, quad and general polygonal patterns. In each such scheme, the subdivision creates a hierarchy of meshes, from the coarse level $\mathcal{M}^0 = (\mathcal{V}^0, \mathcal{E}^0, \mathcal{F}^0)$, which we get as input, to some fine level $\mathcal{M}^k = (\mathcal{V}^k, \mathcal{E}^k, \mathcal{F}^k)$. The vertices of level $k + 1$ comprise new edge or face points (or both), and the original vertices of level k are either the same in level $k + 1$ in interpolatory schemes, or averaged in approximating schemes. To fit our methodology, we categorize the linear subdivision schemes we use into two groups:

- **1-ring subdivision schemes**, where the stencil of every point at level $k + 1$ is contained in one or more 1-rings completely.
- **2-ring subdivision schemes**, where the stencil is not contained in any single 1-ring, but rather contained in the union of two adjacent 2-rings.

The purpose of this categorization is as follows: in the second step of our algorithm, we compute candidate points within each ring individually. In the fourth step, we blend the candidate points contributed from the different rings into the final new vertex positions.

Due to this, for 1-ring schemes, we blend points that are computed using the same stencil, but through different canonical forms. In 2-ring schemes, we need to break the stencils to partial contributions from the independent rings. We next describe these procedures.

1-ring subdivision schemes. A linear stationary subdivision scheme is defined by a constant local (averaging) matrix S^k that describes the positions of new face, edge, and vertex points at level $k + 1$, from the vertices of level k . Each row $S^k(p, \cdot)$ encodes the stencil of the subdivision scheme for a single vertex p by its non-zero elements $\text{NZ}(S^k(p, \cdot))$. Then, we define a subdivision scheme to belong to the “1-ring” categorization when for every 1-ring R_v^k of vertex v at level k , and every stencil $\text{NZ}(S^k(p, \cdot))$ for any new point p at level $k + 1$, we have either $\text{NZ}(S^k(p, \cdot)) \subset R_v^k$ or $\text{NZ}(S^k(p, \cdot)) \cap R_v^k = \emptyset$. In words, the stencil of each point in level $k + 1$ is contained in one

or more 1-rings completely. As such, when using the original linear schemes (without our proposed Möbius modification steps), all the candidate points are trivially equal, and do not require blending. Examples of stencils that exhibit this property:

- (i) A face point as an average of the vertices adjacent to the face.
- (ii) An edge point as an average of the vertices adjacent to the edge.
- (iii) Vertex points that are the weighted average of such face and edge points, and other vertices in their 1-ring.
- (iv) Boundary vertices and edges blended from adjacent boundary vertices alone.

Most schemes, especially approximative schemes, belong to this category: for instance, Loop, Catmull-Clark, corner-cutting and Simplest subdivision [Peters and Reif 1997]. To these we include subdivision-like operators like dual truncation, vertex insertion, and dual meshes, presented in [Akleman et al. 2005].

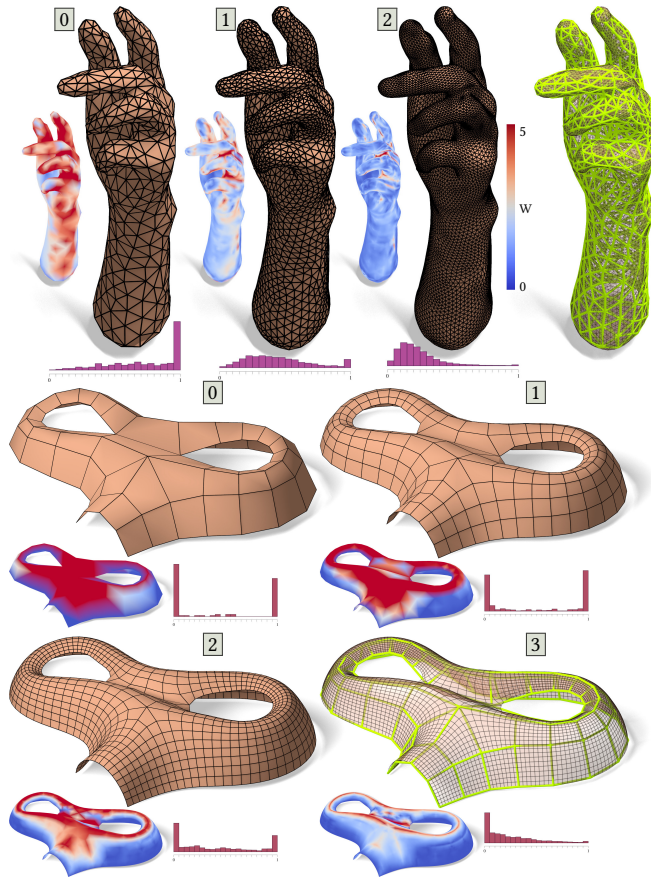


Fig. 9. Triangle and quad pattern subdivision with our method. The color figures show the Willmore energy (marked “W”) of every vertex, measured by the planarity of the tangent polygon. The planarity itself is measured as the RMSE of the planarity values of each four consecutive quads in the tangent polygon. Individual quad planarity is measured in turn by the percentage of the distance between the diagonals to the average diagonal length. Note how the Willmore energy reduces with each refinement level.

2-ring subdivision schemes. The stencils of these schemes are wider and are not contained in any single 1-ring. This is usually a property of interpolatory schemes that require bigger stencils to achieve smoothness. Specifically, we use Kobbelt quadrilateral scheme [Kobbelt 1996] and the Modified Butterfly scheme [Dyn et al. 1990; Ling et al. 2006; Zorin et al. 1996] that exhibit this property. Our modus operandi is to create “partial candidate points”. That is, we break each stencil into the individual contributions from each ring, compute them in their canonical form, and then blend them on the original mesh. This is similar to the process for 1-ring schemes, except that the stencil needs to be adapted. We do that for the Kobbelt scheme and the Modified Butterfly scheme as follows:




Fig. 10. Our subdivision with the Butterfly scheme (the Epcot mesh, *top*), and the Kobbelt scheme (Duck mesh, *bottom*). Note that for the Epcot mesh, the small pyramids become spherical pockets, and that the dimple in the Duck’s nose remains intact, rather than averaged out.

Kobbelt scheme. A regular interior vertex in a quad mesh has valence 4 whereas singularities in the interior are characterized by a valence different from 4. On the boundary we have valences of 2 for corners, 3 for regular boundary vertices or any other number for singular boundary vertices. The Kobbelt scheme creates an edge point v_{jk} with stencils comprising vertices v_j and v_k , respectively, and with given respective virtual points v_i and v_l (see Figure 11 left). For regular vertices v_j and v_k , the virtual points are just the next points along the parameter lines; otherwise, we compute them with the Kobbelt linear stencil in the respective canonical forms C_j and C_k .

The original linear Kobbelt scheme provides linear blending rules for the edge points from the four points v_i, v_j, v_k, v_l , and for the face points as tensor products. We, on the other hand, blend the four points to an edge points in a Möbius invariant way (with the operator $F4$), and blend the face points with the operator $F6$, as explained in Section 4.4.




Fig. 11. Creating partial candidate edge points in 2-ring schemes. Left: In the Kobbelt scheme we create virtual points v_j and v_l in the respective canonical forms, and then blend them to v_{jk} using $F4$. Right: edge points in the Butterfly scheme are created by breaking the 2-ring support into individual 1-ring stencils, creating partial candidate edge points $v_{e,j}$ and $v_{e,k}$, computed on the canonical form. These points are then blended with $F4$ to form v_{jk} .

Modified Butterfly scheme. The edge points of the Butterfly scheme are computed by a symmetric stencil within two neighboring 1-rings, as depicted in Figure 11 (right). Consider edge $e = (v_j, v_k)$ on which we would like to compute edge point v_{jk} . In the case where one of the rings is singular and the other one is not, new points are computed only with regards to the singular ring, and it is like the 1-ring scheme. In case both rings are either singular or regular together, and as the scheme is symmetric, we first compute the partial candidate points $v_{e,j}$ and $v_{e,k}$ as follows. Write the vertices of the first 1-ring in a vector I_j , and of the second in I_k (the two vectors contain mutual vertices as the 1-rings are overlapping). Then, the new edge point v_{jk} can be written as

$$v_{jk} = W_j I_j + W_k I_k = \frac{2W_j I_j + 2W_k I_k}{2},$$

where W_j and W_k are the vectors of vertex weights in the Butterfly scheme. In that formulation the new point v_{jk} is the arithmetic mean of

$$v_{e,j} = 2W_j I_j, \quad \text{and} \quad v_{e,k} = 2W_k I_k,$$

which is just the linear blending of v_{jk} from the partial candidate points $v_{e,j}$ and $v_{e,k}$ of the two involved 1-rings. Our method converts the Butterfly scheme into a Möbius-invariant scheme by computing $v_{e,j}$ and $v_{e,k}$ in the individual canonical forms of j and k , and then blending them in a Möbius invariant way (with the $F4$ -operator), as explained in Section 4.4.

We depict the results of our subdivision for 1-ring schemes in Figures 8 and 9, and for 2-ring schemes in Figure 10.

Boundary curves. In all methods, boundaries (and sharp features) in theory, though we do not enforce them in our examples) are treated as independent curves that are subdivided using curve-based rules. Thus, our Möbius subdivision follows the same paradigm, and uses the (approximating and interpolating) Möbius-invariant curve subdivision rules we present in Section 4.4.1 in lieu of linear curve subdivision rules.

4.4 Möbius-invariant blending

In the fourth and last step of the algorithm, we take the candidate points computed in the linear schemes on the individual canonical forms, after transforming them back to the original mesh, and blend them to the final points of the subdivision. Our blending is based on two novel Möbius-invariant operators that we detail in the following.

4.4.1 Four-point blend $F4$. Consider four points in \mathbb{R}^3 , given as imaginary quaternions: $a, b, c, d \in \text{Im } \mathbb{H}$. We construct a new point p in a Möbius-invariant way: p is defined as the solution to

$$\text{cr}[c, a, b, p] = -\sqrt{\text{cr}[c, a, b, d]}.$$

For the definition of the square root of a quaternion see Sec. 3.1. The solution can be explicitly written as

$$p = ((a-b)(c-a)^{-1}\chi + 1)^{-1} \cdot ((a-b)(c-a)^{-1}\chi c + b), \quad (3)$$

$$\text{where } \chi := -\sqrt{\text{cr}[c, a, b, d]}.$$

This construction leads to a purely imaginary point p , which means a valid point in \mathbb{R}^3 . We provide a proof in Appendix A (Lemma A.1). As the definition of p depends on cross-ratios only, the construction is Möbius-invariant. We denote the resulting point of our four-point insertion rule by $F4(a, b, c, d) := p$. The construction has the following properties:

Cocircularity preserving. Recall that a, b, c, d are cocircular if and only if their cross-ratio $\rho = \text{cr}[c, a, b, d]$ is real. A well-known fact for cross-ratios [Coxeter 1993] is that $\rho > 0$ if and only if we can travel from a to d along the circle without passing through b or c (see Fig. 12 left). As $\sqrt{\rho}$ is positive as well in that case, p must be arranged between b and c . Note that when $b = c$ we get $p = b = c$.

If $\rho < 0$, there are infinitely many quaternionic square roots, corresponding to the case where the pair (b, c) separates the pair (a, d) on the circle like the vertices of the letter “Z” (see Fig. 12 center). In this case, there is no good cocircularity preserving choice; it is instead practical to choose $\chi = \sqrt{|\rho|} \left[0, \frac{(a-b) \times (c-b)}{\|(a-b) \times (c-b)\|} \right]$. Then the new point p does not lie on the same circle but is still in the same plane. We describe this case for completeness, but it is rather contrived and never happens in any of our examples.




Fig. 12. Three examples of the four-point insertion rule. The new point $p = F4(a, b, c, d)$ from concyclic points lies on the same circle iff one can travel from a to d on the circle without meeting either c or b , as in the left, rather than the center case.

Cospherality preserving. Next assume that the four points a, b, c, d are not cocircular, i.e., assuming general position. However, the direction of the imaginary part of $\text{cr}[c, a, b, p]$ is the same as that of χ (again since it is the square root). Consequently, Lemma A.1 immediately implies that the new point p lies on the circumsphere of a, b, c, d (or on the same plane if the four points are coplanar).


Symmetry. Since the construction $F4$ only involves cross-ratios of cospherical points, to compute $F4(a, b, c, d)$ we could also apply a Möbius transformation that maps the circumsphere of a, b, c, d to the complex plane and make our computations there. The tedious part here would be determining the right Möbius transformation. However, this idea gives us the important insight of the symmetry of $F4$, i.e., $F4(a, b, c, d) = F4(d, c, b, a)$, as the complex cross-ratio obeys $\text{cr}[a, b, c, d] = \text{cr}[d, c, b, a]$. Practically, that means it doesn't matter from which direction we blend the points.

Summarizing the above properties we get:

LEMMA 4.1. *The four-point insertion rule $F4(a, b, c, d)$ is circle preserving (iff the pair (a, d) is not separated by (c, b)), sphere preserving, and symmetric, i.e., $F4(a, b, c, d) = F4(d, c, b, a)$.*

Curve subdivision. A beneficial application of $F4$ is the ability to subdivide curves in space with exact sphere and circle reproduction. One can build an interpolating scheme by repeatedly applying $F4$ to every consecutive set of four vertices. It is also possible to build an approximative scheme by further applying $F4$ to the even vertices as well (with the two original even neighbors and the two new odd edge points). We show two examples in Figure 13, and use these operators for boundary-curve subdivision in all our schemes.

4.4.2 Six-point blend $F6$. We consider six points $a, b, c, d, e, f \in \text{Im } \mathbb{H}$ partitioned into two chains: a, b, c as one chain, and d, e, f as another (see inset). We create a point p that interpolates between them. To this end, we use the following construction where we solve for p in $\text{cr}[e, a, b, p] = \chi$, with χ defined as below. In a similar manner to $F4$, we can write p explicitly as



$$p = ((a - b)(e - a)^{-1} \chi + 1)^{-1} \cdot ((a - b)(e - a)^{-1} \chi e + b),$$

$$\text{where } \chi := -\sqrt{\text{cr}[e, a, b, d]} \cdot \sqrt{r},$$

$$\text{with } r := \sqrt{\text{cr}[e, a, b, d]}^{-1} \cdot \text{cr}[e, a, b, f] \cdot \sqrt{\text{cr}[e, c, b, f]}^{-1}.$$

While this construction of p seems convoluted, note that all cross-ratios begin with e : in the case where all points are on the same sphere, we then have a blend of cross-ratios with co-directional




Fig. 13. Our algorithm for point insertion with $p = F4(a, b, c, d)$ results in circle- and sphere-reproducing curves. *Top:* interpolatory circular curve. *Bottom:* approximating spherical curve. Control polygon in brown.

imaginary parts (since the inverse, square root, and product of quaternions do not change this direction). p , constructed this way, is always imaginary. For a proof again apply Lemma A.1 (see Appendix A). In the case where the six points are not cospherical we observed that p is still imaginary in all present cases of our examples and in numerical experiments. However, we do not yet have a rigorous proof for that case. We denote this construction by $F6(a, b, c, d, e, f) := p$ and for the same reason as in the four-point insertion rule and as detailed in Appendix A we obtain:

LEMMA 4.2. *For six cospherical points the six-point insertion rule $F6(a, b, c, d, e, f)$ is sphere preserving, i.e., the new point lies on the same sphere. If $b = e$, we get $p = b = e$.*

4.4.3 Blending edge and face points. As the final step of our algorithm we apply the blending operators $F4$ and $F6$ to candidate edge and face points computed in the canonical forms and transformed back to the original mesh. We explain the rules for blending edge and face points for 1-ring schemes as follows.

Blending edge points. In 1-ring schemes, for each two neighboring 1-rings with central vertices v_j and v_k , and the edge e between them, we obtained two candidate edge points $v_{e,j}$ and $v_{e,k}$ computed from the two canonical forms and transformed back. The final edge point is computed as:

$$v_{jk} = F4(v_j, v_{e,j}, v_{e,k}, v_k).$$

Note that if the mesh is Möbius regular, it is Möbius equivalent to a regular grid (at least locally). In that case, we get $v_{e,j} = v_{e,k} = v_{jk}$, where in both canonical forms this is merely the mid-edge point in all the subdivision schemes we use. For instance, a perfect grid is subdivided into a perfect grid in Catmull-Clark subdivision (see Figure 15), and our scheme reproduces this naturally under a Möbius transformation.

In both the Kobbelt and the Butterfly scheme (Figure 11) we have a similar construction for edge points: for the Kobbelt scheme, we blend two edge points v_j and v_k with the virtual points v_i and v_l as $v_{jk} = F4(v_i, v_j, v_k, v_l)$, and in the Butterfly scheme we have the same construction as in 1-ring methods, except that the candidate

edge points are not equal in general, as they are partial candidate points computed each with a parted linear stencil.




Fig. 14. Computing the face point in the Kobbelt scheme. Having computed edge points $w_{1\dots 8}$ (yellow), we compute candidate face points $w_{1234} = F4(w_1, w_2, w_3, w_4)$ and $w_{5678} = F4(w_5, w_6, w_7, w_8)$ (light blue). Then we compute the final face point $p = F6(w_2, w_{1234}, w_3, w_6, w_{5678}, w_7)$.

Blending face points. In the linear Kobbelt scheme, the virtual points are designed in such a way that the face point can be computed in either direction of the quad mesh, horizontally or vertically (i.e., as tensor product of the 4-point scheme). Unfortunately, this is not guaranteed in our scheme unless the mesh is perfectly Möbius regular; as such, for any face, we compute two candidate face points, one blended with the vertical edge points, and one with the horizontal edge points, and compute them as illustrated in Figure 14.

Next consider face points for faces of valence d in 1-ring schemes. We get a candidate face point $v_{f,i}$ from transforming back the candidate points computed in the canonical forms C_i of the 1-rings around each vertex v_i . We perform the following algorithm to compute the final single new face point v_f :

- (i) Pair every vertex v_i with opposite vertex $v_{i+\frac{d}{2}}$. Create set A of quadruplets $v_i, v_f, v_{f,i+\frac{d}{2}}, v_{i+\frac{d}{2}}$.
- (ii) Sort A according to the ascending noncircularity or the quadruplets, measured as $1 - |\text{Re}(\text{cr})|/|\text{cr}|$ (difference of cross ratio of quadruplet from a real number).
- (iii) For each quadruplet $(a_l, b_l, c_l, d_l) \in A$, compute the new point $p_l = F4(a_l, b_l, c_l, d_l)$.
- (iv) Assign $v_f = p_1$, where p_1 is the p of the first quadruplet in A .
- (v) For $l = 2, \dots, d$ compute $v_f := F6(a_{l-1}, v_f, c_{l-1}, a_l, p_l, c_l)$.

We use all indices modulo d , and if d is odd, we round down $\frac{d}{2}$. In words, the algorithm blends the candidate face points one by one, where the ones that are already on a circle blend before others, as they are less “noisy”. Note that the sorting is cross-ratio based and thus doesn’t break the Möbius invariance. However, the algorithm is not symmetric with regards to the points in the face; we found that this has little impact in practice. Note that in this case as well, if the 1-rings around the mesh are all Möbius regular, all candidate points are equal to the final result $v_{f,i} = v_f$. As the face is Möbius-regular in that case as well, consider the regular n -gon to which the face is

Möbius-equivalent, and its center c ; then, v_f is the image of Möbius transformation of that center.

4.5 Properties of canonical Möbius subdivision

The stencil of new edge and face points, and deformed vertex points, is limited to a 2-ring environment of the coarsest control cage. As such, we can provide some theoretical guarantees on spherical and regularity reproduction of elements in the refined mesh.

Spherical reproduction. If a vertex of level k and the support of its stencil are on a generalized sphere (i.e., sphere or plane), then it is on the same sphere for all levels $m \geq k$. To see that, take a Möbius transformation that takes the entire support to a plane, which is possible due to cosphericity. Then, every linear subdivision, sharing the same support, would preserve the plane, and the Möbius transformation back would reproduce the original sphere. As an interesting consequence, even when we use approximative subdivisions, locally cospherical environments often interpolate or nearly interpolate the coarse points, just because the sphere is reproduced. See such examples in Figures 8, 16, and 18.

Möbius-regularity reproduction. By a similar logic, if the support of a vertex at level k is Möbius equivalent to a Euclidean regular pattern in the plane (see Fig. 15 bottom), and the chosen underlying linear scheme preserves Euclidean regularity, then the vertex is Möbius-regular for all levels $m \geq k$. Note that this a stronger requirement than just having every vertex in the support be Möbius regular on its own, as they do not necessarily have a mutual single Möbius transformation to the plane. See Fig. 15 top for a counterexample.




Fig. 15. Möbius regularity under Möbius subdiv. *Top*: the octahedron is Möbius regular, and the stencil for vertex points is as well, but not the stencil for edge points, and thus the subdivision loses regularity at the latter. Note that the irregularity diffuses under refinement. MR measures the squared sum of differences between cross-ratios on each tangent polygon ($0 =$ Möbius regular). *Bottom*: the Möbius subdiv. of a perfect grid under a Möbius transformation preserves the regularity perfectly.




Fig. 16. The impulse responses for Catmull-Clark (left) and Loop (right) Möbius subdivision.

Bijectivity. Consider a (generalized) spherical 1-ring in a mesh with central vertex v . We say it is non-flipping if the boundary polygon B_v is a non-self-intersecting curve that divides the sphere into “outside” and “inside” regions by the orientation of B_v , and where v is inside. A subdivision is *bijjective* if it preserves this property under refinement. By empirical evidence on all our examples, we conjecture that if a vertex is *Delaunay*, which is defined as having a convex tangent polygon, and so are all its neighbors, then the subdivision is locally bijective. We cannot provide a formal proof and leave this to future work; nevertheless, the intuition behind this is that the chosen center $c(T_v)$ for the $m_{B_v \rightarrow C_v}$ is inside the tangent polygon, and therefore the transformation does not have a pole inside T_v , and thus avoids flipping triangles. See further evidence in Sections 6.4 and 7.

5 ANALYSIS

Impulse response. We show the “impulse response” (the non-linear analogue of a basis function) for the Catmull-Clark and Loop Möbius subdivisions in Figure 16. They are made by subdividing a perfect Euclidean regular grid at $z = 0$ with one vertex at $z = 1$. This highlights the spherical nature of the subdivision. Note that Loop is interpolating in this case; that is because the created cone is cospherical.

Spherical reproduction. We subdivide several distorted and uneven, yet perfectly spherical meshes with both Möbius subdivision and linear subdivision in Figure 17. This demonstrates our guaranteed property of perfect spherical reproduction. We also demonstrate it in the same figure on polyhedral pattern subdivision.

Linear and Möbius subdivisions. In Figure 18, we analyze the behaviour of Möbius subdivision with relation to benchmark subdivision meshes. The star mesh is subdivided by Butterfly and Loop schemes, both linear and Möbius. The star has cones (spikes) with low Willmore energy, and adjacent hyperbolic junctions between them that are far from spherical (see Willmore energy plots). As a consequence, Möbius subdivision tends to turn these cones into approximate parts of spheres considerably more than the linear subdivisions. Another consequence is that Möbius Loop subdivision almost interpolates the cones, while the linear subdivision naturally shrinks the features as a result of averaging.

In the case of the T mesh, we see a similar behavior: the corners of the T are cospherical with their 1-rings (but their neighbors are not), and therefore Möbius Catmull-Clark subdivision almost interpolates these nodes. Möbius subdivision creates meshes that




Fig. 17. Sphere reproduction. *Top left* is subdivided using Möbius and linear Kobbelt subdivision. *Top right* using Butterfly. *Bottom row* is subdivided using dual truncation, Loop, dualization, and Catmull-Clark. All results have vertices with perfect norm 1 on the unit sphere.




Fig. 18. *Top left*: Butterfly linear and Möbius Butterfly subdivision of the star. *Top right*: the same for Loop subdivision. *Middle row*: Catmull-Clark linear and Möbius subdivisions. *Bottom row*: Kobbelt subdivision.

look like a blend of spheres for this mesh, while the linear Catmull-Clark shrinks as the result of averaging, and the Kobbelt subdivision creates sharp corners.

Nevertheless, while our algorithm is as parallel and computationally efficient in the asymptotic sense as the linear method it modifies, it bears some overhead locally. This is because linear methods only

Table 1. Time overhead of Möbius subdivision vs. linear subdivision stress test. The numbers # coarse and # fine indicate face counts, and the times “linear” and “Möbius” are in seconds, measuring total subdivision time from coarse to fine. “Ratio” is the overhead of Möbius vs. linear. Note that quad mesh methods are in general more expensive because of the computation of face points.

	Duck (Kobbelt)	Epcot (Butterfly)	Star (Loop)	T-mesh (CC)
# coarse	334	1536	24	18
# fine	342016	393216	393216	294912
Linear	0.48	1.4	0.22	0.32
Möbius	1.59	3.46	0.59	1.8
Ratio	3.3	2.47	2.68	5.6

Table 2. Maximal edge length as a function of iteration. The convergence rate is measured as 10^{-s} , where s is the regression slope of the logarithm of the data (larger s means more contractive).

Level	Duck (Kobbelt)	Epcot (Butterfly)	Star (Loop)	T-mesh (CC)
0	0.305	0.303	3.317	2.062
1	0.189	0.153	2.049	1.296
2	0.101	0.089	1.204	0.688
3	0.053	0.049	0.634	0.350
4	0.027	0.025	0.327	0.176
Rate	1.846	1.847	1.788	1.865

require one matrix multiplication, and our method requires computing canonical forms, and several expensive quaternionic operations. We show the difference in overhead by constructing a “stress test” that measures the time of subdividing a mesh to a very fine level in several methods, and show the results in Table 1. While the overhead seems substantial, a vectorization of both methods would render the total time difference negligible for any practical purpose.

Empirical convergence. While we don’t provide any formal convergence proof for our subdivision scheme (except in the cases of Möbius regularity and perfect spheres), in Table 2 we show that our scheme is empirically *contractive* in terms of maximal edge length, for the same examples as in the stress test. It is evident that there is a consistent near-quadratic order of convergence.

6 APPLICATIONS

In the following, we show how Möbius subdivision and its properties are useful for several existing applications in geometry processing.

6.1 Implementation Details

Our code is run on a 4GHz i7 iMac with 32GB memory. While our subdivision overhead is easily parallelizable, in our implementation the subdivision is run sequentially on each vertex in each level and the code is therefore not optimized. However, the local nature of the subdivision allows for even more substantial speedup with a future parallel implementation on the GPU. We used libhedra [Vaxman 2016] for Möbius geometry on polygonal meshes. The spherical

orbifold optimization is run with code provided openly by the authors [Aigerman et al. 2017], and our curve subdivision (Figure 13) is coded in MATLAB.

6.2 Coarse-to-fine editing

Coarse-to-fine mesh editing proceeds by editing meshes in increasing order of refinement, with the intend of adding coarse details first, and fine details upon them. The editing is performed on the chosen level-of-detail coarse cage, interactively subdivided into the finest level. Our algorithm easily lends to coarse-to-fine mesh editing by subdivision surfaces: we pick and drag vertices of a control polygon of level k , and alter all subsequent levels by Möbius subdivision. The tendency of Möbius subdivision to result in smooth and as-spherical-as-possible meshes is demonstrated in Figure 19. Our subdivision does not add significant overhead to the linear subdivision schemes, and therefore this editing was interactive (≈ 1 second lag for updating coarse-to-fine). See video for interaction.

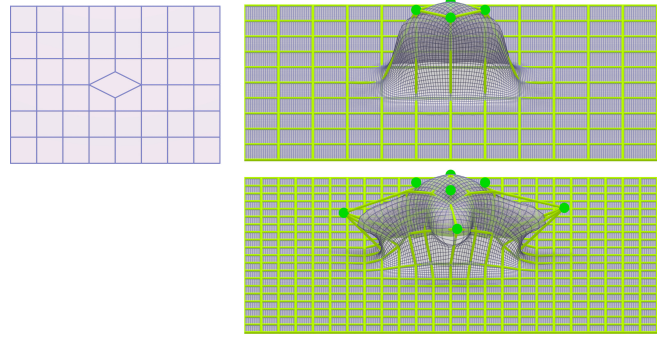


Fig. 19. Coarse-to-fine editing. *Left*: original pattern in top view. *Top right*: editing in level 0. *Bottom right*: further editing in level 1. The final mesh is at level 4. The editing is done with Catmull-Clark Möbius subdivision, and therefore the handles are in general not interpolating.

6.3 Efficient regular meshes

In addition to directly editing meshes, we can make the optimization of [Vaxman et al. 2017] considerably more efficient by optimizing a coarse level and subdividing to a fine level, as our subdivision approximately preserves as-Möbius-regular-as-possible meshes. We demonstrate the order of magnitude scale in efficiency in Figure 20, which makes the costly nonlinear optimization of [Vaxman et al. 2017] reach interactive rates, without considerable difference in the result. See video for an example as well.

6.4 Spherical Orbifolds

Aigerman et al. [2017] present a spherical parameterization method that maps a closed genus-0 mesh into a patch that tiles the sphere, as a generator of an orbifold structure. The condition for perfect tiling is enforced by having several boundaries be rigid transformations of each other. The algorithm uses L-BFGS nonlinear optimization, and therefore does not scale well to large meshes. Möbius subdivision is useful for efficient spherical orbifold computation for the following properties: first, spherical reproduction. Second, if a vertex star

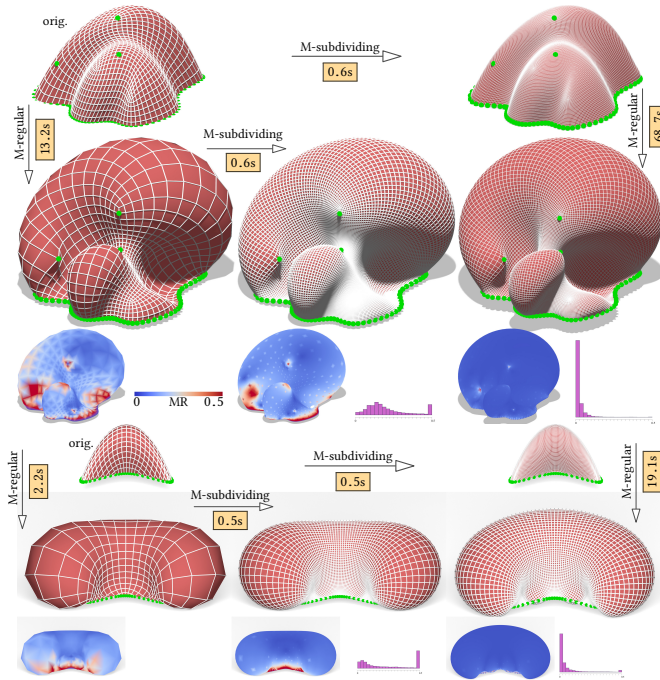


Fig. 20. Efficient regular meshes. *Top*: coarse bubble mesh (1076 faces), subdivided twice (for 17216) meshes. The coarse optimization with subdivision takes $13.2 + 0.6 = 13.8$ seconds, while direct optimization on the subdivided mesh takes 68.7 seconds. The fine result is numerically more regular, but the difference is mild. *Bottom*: metro mesh (coarse: 464 faces, fine: 7424 faces), where optimization and the subdivision takes $2.2 + 0.5 = 2.7$ seconds, while fine-level optimization takes 19.1 seconds. Here again the difference is apparent, but small.

has rotational symmetry, then by construction the canonical form also maintains the same symmetry. As such, if we “glue” the tiles together, subdivide the resulting spherical mesh, and segment them again, the resulting subdivided patches are still an orbifold tiling. We demonstrate this coarse-to-fine spherical orbifolds in Figure 21. Note that as the original meshes are Delaunay, the result also preserves bijectivity in the demonstrated examples; however, we cannot give a theoretical guarantee that bijectivity is preserved as a rule (see Section 4.5).

Since we commute with Möbius transformations, it is worth exploring whether we can create subdivision schemes that truly reproduces conformal (or harmonic) transformations, and we leave this for future work.

7 DISCUSSION

7.1 Limitations of canonical subdivision

Accentuating blobs. Spherical and regularity reproduction are appealing properties in several contexts. Nevertheless, it might not always be the case. Consider the Monster Frog case in Figure 22. The “wavy” stomach is averaged out in linear subdivision schemes, but our method accentuates the wave into blobs. Therefore, Möbius subdivision is only good when such a result is expected.




Fig. 21. Subdividing spherical orbifolds. Each vertical half of the image is for a different mesh. *First row, left to right*: coarse original mesh, subdivided original mesh, coarse-computed orbifold, fine-computed orbifold (same cones, but different cuts, to which the method is invariant), subdivided coarse-computed orbifold. *Third row*: corresponding normal maps encoded in RGB space for subdivided coarse-computer and fine-computed. Zoom-ins are provided for all examples above them, focusing on a single cone. Computed times are reported on the image. It is evident that coarse orbifold + subdivision scale much better than direct fine orbifold optimization, with similar results, and that the seamless continuity is maintained.

Non-Delaunay meshes. As mentioned in Section 4.5, we conjecture that the transformation $m_{B \rightarrow C}$ is not guaranteed to result in bijective refinement when the vertex star is non-Delaunay. In such cases, our algorithm can create artifacts in the subdivision in some extreme cases. We exemplify this in Figure 23, where some 1-rings in the Bimba mesh fold over in Möbius subdivision. Remeshing then alleviates this problem.

7.2 Future Work

It is possible to use our subdivision as a full-blown multiresolution framework, by defining detail functions encoded linearly in the canonical forms. In addition, it is interesting to see if Möbius




Fig. 22. Blobby behaviour. The wavy stomach of the Monster Frog (left) is averaged out in linear subdivision (center), but accentuated in Möbius subdivision (right).




Fig. 23. non-Delaunay subdivision can create artifacts in some vertex stars. Left to right: original Non-Delaunay Bimba, subdivided mesh with flipped lower-left shoulder, remeshing, and then subdivision which preserves the shape well. Here, we show the Euclidean regularity of the faces as measured by [Vaxman et al. 2017]: RMSE of sum of squared differences of the quaternionic ratios of adjacent edges in any single face.

subdivisions can be blended with linear subdivisions as a spectrum, by interpolating the canonical forms with the original vertex stars. We will explore such extensions in future work.

We consider as future work rigorous proofs of bijectivity properties, as evidenced in Section 4.5. In addition, it is necessary to prove complete theorems about smoothness and convergence of the scheme. We did not find any examples where the scheme did not converge in practice, and conjecture that it is because of the conversion of the existing well-known linear schemes, as done in other nonlinear methods.

Finally, it is interesting to explore subdivision schemes that subdivide two conformally-equivalent meshes simultaneously, in a way that preserves their conformal equivalence, under some definition of discrete conformality that is Möbius related. We expect that our scheme can inspire such a subdivision method as well.

ACKNOWLEDGMENTS

This research was partially funded by the Austrian Science Fund (FWF) (projects P 29981 and I 2978-N35) and by the Israel Science Foundation (grants No. 1869/15 and 2102/15). We thank Ron Vanderfeesten for the stylistic rendering of Figure 2.

A POINTS ON SPHERES

LEMMA A.1. *Let $a, b, c, d \in \text{Im } \mathbb{H} \cong \mathbb{R}^3$ be four points in general position with $\text{cr}[a, b, c, d] = [r, v]$. Further, let $p \in \mathbb{H}$ be the quaternion that solves $\text{cr}[a, b, c, p] = [\lambda r, \mu v]$, for some $\lambda, \mu \in \mathbb{R}$ (the explicit formula for p as solution of the above equation is given by Equation (3) where $\chi = [\lambda r, \mu v]$). Then $p \in \text{Im } \mathbb{H}$, i.e., p is an imaginary quaternion. Furthermore, p lies on the circumsphere of a, b, c, d . In particular $p(\lambda, \mu)$ is a parameterization of the circumsphere.*

PROOF. From Equation (2), we know that the two occurring cross-ratios can be expressed in terms of corner tangents as

$$\text{cr}[a, b, c, d] = t_1 \cdot t_2, \quad \text{and} \quad \text{cr}[a, b, c, p] = t_1 \cdot t_3,$$

where we denote for brevity $t_1 := t[c, a, b]^{-1}$, $t_2 := t[c, a, d]$, and $t_3 := t[c, a, p]$. Consequently, as all $t_i \in \text{Im } \mathbb{H}$, we have

$$[r, v] = [-\langle t_1, t_2 \rangle, t_1 \times t_2], \quad \text{and} \quad [\lambda r, \mu v] = [-\langle t_1, t_3 \rangle, t_1 \times t_3].$$

Since the corner tangents t_1, t_2, t_3 are orthogonal to v and therefore in a plane we can express t_3 in the form $t_3 = \alpha t_1 + \beta t_2$. Together with the above equation we obtain

$$\begin{aligned} \lambda \langle t_1, t_2 \rangle &= -\lambda r = \langle t_1, t_3 \rangle = \alpha \langle t_1, t_1 \rangle + \beta \langle t_1, t_2 \rangle, \\ \mu t_1 \times t_2 &= \mu v = t_1 \times t_3 = \alpha t_1 \times t_1 + \beta t_1 \times t_2. \end{aligned}$$

Consequently, $\beta = \mu$ and $\alpha = (\lambda - \mu) \langle t_1, t_2 \rangle / |t_1|^2$, which determines t_3 uniquely. From the definition of $t_3 = t[c, a, p] = (c - a)^{-1} (p - a)^{-1}$, we then immediately get

$$p = (t_3 - (a - c)^{-1})^{-1} + a \in \text{Im } \mathbb{H}.$$

Furthermore, the circumsphere of a, b, c, p is the same as the circumsphere of a, b, c, d since both pass through a, b, c and both have parallel normal vecctors (μv and v , resp.) at a , and there is only one such sphere. \square

REFERENCES

- Aigerman Noam, Kovalsky Shahar Z., and Lipman Yaron. 2017. Spherical Orbifold Tutte Embeddings. *ACM Trans. Graph.* 36, 4 (2017), 90:1–90:13.
- Aklemann Ergun, Srinivasan Vinod, and Mandal Esan. 2005. Remeshing Schemes for Semi-Regular Tilings. In *Proc. SMI*. IEEE Computer Society, 44–50.
- Bo Pengbo, Pottmann Helmut, Kilian Martin, Wang Wenping, and Wallner Johannes. 2011. Circular Arc Structures. *ACM Trans. Graph.* 30, 4 (2011), 101:1–101:12.
- Bobenko Alexander and Pinkall Ulrich. 1996. Discrete isothermal surfaces. *Journal for Pure and Applied Mathematics* 475 (1996), 187–208.
- Bobenko Alexander I. and Schröder Peter. 2005. Discrete Willmore Flow. In *Proceedings of the Third Eurographics Symposium on Geometry Processing (SGP '05)*. Eurographics Association, 101–110.
- Bobenko Alexander I. and Springborn Boris A. 2004. Variational principles for circle patterns and Koebe's theorem. *Trans. Amer. Math. Soc.* 356, 2 (2004), 659–689.
- Bouaziz Sofien, Deuss Mario, Schwartzburg Yuliy, Weise Thibaut, and Pauly Mark. 2012. Shape-Up: Shaping Discrete Geometry with Projections. *Comput. Graph. Forum* 31, 5 (2012), 1657–1667.
- Cashman Thomas J. 2012. Beyond Catmull-Clark? A Survey of Advances in Subdivision Surface Methods. *Comput. Graph. Forum* 31, 1 (2012), 42–61.
- Catmull Edwin and Clark Jim. 1978. Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design* 10, 6 (1978), 350–355.
- Cavaretta Alfred S., Micchelli Charles A., and Dahmen Wolfgang. 1991. *Stationary Subdivision*. American Mathematical Society, Boston, MA, USA.

- Coxeter H. S. M. 1993. *The real projective plane* (third ed.). Springer-Verlag, New York. xiv+222 pages.
- Crane Keenan, Pinkall Ulrich, and Schröder Peter. 2013. Robust Fairing via Conformal Curvature Flow. *ACM Trans. Graph.* 32, 4 (2013), 61:1–61:10.
- Dyn Nira, Levine David, and Gregory John A. 1990. A Butterfly Subdivision Scheme for Surface Interpolation with Tension Control. *ACM Trans. Graph.* 9, 2 (1990), 160–169.
- Floater Michael and Micchelli Charles A. 1998. Nonlinear stationary subdivision. *Approximation Theory: In Memory of A.K. Varma* 86, 1 (1998), 209–224.
- Jiang Caigui, Tang Chengcheng, Seidel Hans-Peter, and Wonka Peter. 2017. Design and Volume Optimization of Space Structures. *ACM Trans. Graph.* 36, 4 (2017), 159:1–159:14.
- Jiang Caigui, Tang Chengcheng, Vaxman Amir, Wonka Peter, and Pottmann Helmut. 2015. Polyhedral Patterns. *ACM Trans. Graph.* 34, 6 (2015), 172:1–172:12.
- Kharevych Liliya, Springborn Boris, and Schröder Peter. 2006. Discrete Conformal Mappings via Circle Patterns. *ACM Trans. Graph.* 25, 2 (2006), 412–438.
- Kobbelt Leif. 1996. Interpolatory Subdivision on Open Quadrilateral Nets with Arbitrary Topology. *Comput. Graph. Forum* 15, 3 (1996), 409–420.
- Ling Ruotian, Luo Xiaonan, Chen Ren, and Zheng Guifeng. 2006. An Interpolatory Subdivision Scheme for Triangular Meshes and Progressive Transmission. In *Interactive Technologies and Sociotechnical Systems*, Hongbin Zha et al. (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 242–252.
- Loop Charles Teorell. 1987. Smooth subdivision surfaces based on triangles. Master's thesis, University of Utah, Department of Mathematics.
- Micchelli Charles A. 1996. Interpolatory Subdivision Schemes and Wavelets. *Journal of Approximation Theory* 86, 1 (1996), 41–71.
- Peng Chi-Han, Pottmann Helmut, and Wonka Peter. 2018. Designing Patterns Using Triangle-quad Hybrid Meshes. *ACM Trans. Graph.* 37, 4, Article 107 (July 2018), 14 pages. <https://doi.org/10.1145/3197517.3201306>
- Peters Jörg and Reif Ulrich. 1997. The Simplest Subdivision Scheme for Smoothing Polyhedra. *ACM Trans. Graph.* 16, 4 (1997), 420–431.
- Pottmann Helmut, Eigensatz Michael, Vaxman Amir, and Wallner Johannes. 2015. Architectural Geometry. *Computers and Graphics* 47 (2015), 145–164.
- Pottmann Helmut, Liu Yang, Wallner Johannes, Bobenko Alexander, and Wang Wenping. 2007. Geometry of Multi-layer Freeform Structures for Architecture. *ACM Trans. Graph.* 26, 3 (2007), 65:1–65:11.
- Rahman Inam Ur, Drori Iddo, Stobben Victoria, Donoho David.L., and Schröder Peter. 2005. Multiscale Representations for Manifold-Valued Data. *Multiscale Modeling & Simulation* 4, 4 (2005), 1201–1232.
- Sabin Malcolm A. and Dodgson, Neil A. 2005. A circle-preserving variant of the four-point subdivision scheme. In *Mathematical Methods for Curves and Surfaces: Tromsø 2004, Modern Methods in Mathematics*. Nashboro Press, 275–286.
- Schaefer Scott, Vouga Etienne, and Goldman Ron. 2008. Nonlinear subdivision through nonlinear averaging. *Computer Aided Geometric Design* 25, 3 (2008), 162–180.
- Springborn Boris, Schröder Peter, and Pinkall Ulrich. 2008. Conformal Equivalence of Triangle Meshes. *ACM Trans. Graph.* 27, 3 (2008), 77:1–77:11.
- Tang Chengcheng, Sun Xiang, Gomes Alexandra, Wallner Johannes, and Pottmann Helmut. 2014. Form-finding with Polyhedral Meshes Made Simple. *ACM Trans. Graph.* 33, 4 (2014), 70:1–70:9.
- Vaxman Amir. 2016. libhedra: geometric processing and optimization of polygonal meshes. <https://github.com/avaxman/libhedra>.
- Vaxman Amir, Müller Christian, and Weber Ofir. 2015. Conformal mesh deformations with Möbius transformations. *ACM Trans. Graph.* 34, 4 (2015), 55:1–55:11.
- Vaxman Amir, Müller Christian, and Weber Ofir. 2017. Regular Meshes from Polygonal Patterns. *ACM Trans. Graphics* 36, 4 (2017), 113:1–113:15.
- Wallner Johannes and Dyn Nira. 2005. Convergence and C1 Analysis of Subdivision Schemes on Manifolds by Proximity. *Comput. Aided Geom. Des.* 22, 7 (2005), 593–622.
- Zorin Denis, Schröder Peter, and Sweldens Wim. 1996. Interpolating Subdivision for Meshes with Arbitrary Topology. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96)*. ACM, 189–192.