

Visual Smoothness of Polyhedral Surfaces

DAVIDE PELLIS, TU Wien

MARTIN KILIAN, TU Wien

FELIX DELLINGER, TU Wien

JOHANNES WALLNER, TU Graz

HELMUT POTTMANN, KAUST

Representing smooth geometric shapes by polyhedral meshes can be quite difficult in situations where the variation of edges and face normals is prominently visible. Especially problematic are saddle-shaped areas of the mesh, where typical vertices with six incident edges are ill suited to emulate the more symmetric smooth situation. The importance of a faithful discrete representation is apparent for certain special applications like freeform architecture, but is also relevant for simulation and geometric computing.

In this paper we discuss what exactly is meant by a good representation of saddle points, and how this requirement is stronger than a good approximation of a surface plus its normals. We characterize good saddles in terms of the normal pyramid in a vertex.

We show several ways to design meshes whose normals enjoy small variation (implying good saddle points). For this purpose we define a discrete energy of polyhedral surfaces, which is related to a certain total absolute curvature of smooth surfaces. We discuss the minimizers of both functionals and in particular show that the discrete energy is minimal not for triangle meshes, but for principal quad meshes. We demonstrate our procedures for optimization and interactive design by means of meshes intended for architectural design.

CCS Concepts: • **Computing methodologies** → **Shape modeling**; *Optimization algorithms*.

Additional Key Words and Phrases: Smoothness, polyhedral surfaces, variation of normals, architectural geometry, total absolute curvature

ACM Reference Format:

Davide Pellis, Martin Kilian, Felix Dellinger, Johannes Wallner, and Helmut Pottmann. 2019. Visual Smoothness of Polyhedral Surfaces. *ACM Trans. Graph.* 38, 4, Article 31 (July 2019), 11 pages. <https://doi.org/10.1145/3306346.3322975>

1 INTRODUCTION

The approximation of a smooth surface by a triangle mesh, or more generally, by a polyhedral mesh, is a basic task of geometry processing. It is a bit surprising that the following aspect has not received more attention in the past, namely the proper representation of saddle-shaped surfaces by meshes. Here ‘proper’ means that the shape of the immediate neighbourhood of a vertex in the mesh should resemble the shape of a small piece of smooth surface. This

Authors’ addresses: Davide Pellis, TU Wien; Martin Kilian, TU Wien; Felix Dellinger, TU Wien; Johannes Wallner, TU Graz; Helmut Pottmann, KAUST.

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 the author(s) 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.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0730-0301/2019/7-ART31 \$15.00

<https://doi.org/10.1145/3306346.3322975>





Fig. 1. A reflective surface reveals the deficiencies in mesh fairness which we quantify by means of a discrete fairness energy based on edge lengths and dihedral angles. This energy is minimal for principal quad meshes (right). In this special case both meshes approximate a reference surface of small total variation of the normal vector field, so all visible deficiencies are not caused by the reference shape, but by the way this shape is meshed.

criterion can to some extent be expressed by the requirement that dihedral angles in the mesh are small.

Meshes which are defective in this respect cause not only problems with the computation of discrete differential quantities, but their visual appearance can be very different from the underlying smooth reference surface. This is particularly true for meshes used for freeform architectural skins with reflective materials.

1.1 Overview and contributions

We start in §2 with a discussion of the inevitable deficiencies of discrete surfaces which are meant to approximate smooth surfaces. We identify the shape of the normal pyramid as a key difference between the discrete case and the smooth case. We propose that a mesh faithfully models a smooth surface if, in negatively curved areas, it exhibits so-called *good saddles*. We go on to show that then normal pyramids are free of self-intersections. This is a criterion for visual smoothness of the mesh (in particular it implies small dihedral angles). We introduce a class of highly-constrained *principal-symmetric* meshes which in some sense optimally represent saddle-shaped surface geometry. The different mesh properties we study can be symbolically arranged in the following order,



with the more restrictive properties to the left.

To quantify the visual smoothness of a mesh, in §3 we turn to a mesh energy measuring the total variation of the normal vectors. It essentially is the sum of edge lengths times dihedral angles. A small energy implies that normal pyramids cannot be too badly behaved.

A thorough discussion of properties of this energy reveals that for a given geometric shape, the mesh with smallest energy is no triangle mesh, but a quad mesh aligned with principal curvature lines. We are not able to strictly prove this statement, but we provide strong motivation. This result is fortunate, because principal meshes are relevant for freeform architecture for various reasons. They admit offsets at constant distance, have near-rectangular panels, and possess torsion-free support structures – for an overview see the survey [Pottmann et al. 2015]. It is interesting to know that they have good properties also from the viewpoint of visual appearance.

It turns out that among all meshes approximating a given smooth reference surface, the lowest achievable energy equals a certain curvature measure which we call *total absolute curvature*, see § 4. It has an interpretation as total variation of the surface’s normal vector field. If that quantity is small, the surfaces have especially good visual smoothness when represented by a mesh.

Finally, in § 5 we use a single computational framework to deal with (i) interactive modeling with principal-symmetric meshes, (ii) the optimization of meshes towards lower energy, and (iii) finding surfaces of minimal total absolute curvature for given boundary.

1.2 Previous Work

We do not attempt to review the wealth of literature on fairness of meshes here. Early quantifications of fairness, e.g. the umbrella vector of Kobbelt et al. [1998] are still relevant to us because of their usefulness as regularizers in nonlinear optimization procedures.

There has been systematic work on meshes which approximate smooth surfaces, and criteria which ensure that discrete-differential quantities derived from meshes approximate their smooth counterparts, see [Bauer et al. 2010; Hildebrandt et al. 2006]. This question is also going to be relevant in our work.

Smoothness of polyhedral surfaces in our narrower sense revolves around the behaviour of the *Gauss image* of the mesh, i.e., the face normals. Günther et al. [2017] investigate in detail properties of vertex stars and Gauss images, establishing relations between the local shape of the mesh and the question of self-intersections of the Gauss image. On that basis, Jiang et al. [2016] discuss optimization of meshes towards star-shaped Gauss images (which implies absence of self-intersections). Good representation of saddles is already implicitly present in their work, but is only expressed in terms of relative position of edges to a smooth reference surface.

It is worth noting that the relation between mesh and Gauss image we employ in this paper has a 2-dimensional analogy, namely the relation between a 2D triangulation and a reciprocal-dual mesh. Here Orden et al. [2004] characterize pseudo-triangulations whose reciprocals have no self-intersections.

Imposing conditions on the normal pyramid has been successfully done before: the “hinge condition” of [Stein et al. 2018] corresponds to developability of surfaces.

To achieve fairness (in our special sense) of meshes, we work with a mesh fairness energy. It falls into a broader class of energies of the form $\sum f(\ell_e, \alpha_e)$, where the sum is taken over all edges of a mesh and f is a function of edge length ℓ_e and dihedral angle α_e , see [Tamstorf and Grinspun 2013] for an overview. We employ the energy $\mathcal{E} = 2 \sum \ell_e |\sin \frac{\alpha_e}{2}|$ which approximates the energy




Fig. 2. *Discrepancy between smooth surfaces and discrete surfaces.* (a) When circling a point v of a smooth surface Φ , the normal vector describes a convex cone. (b) In a polyhedral surface, the normal vectors of a vertex star form a pyramid. Deviating from the smooth situation, this pyramid is typically non-convex in case of negative curvature. (c) Normal pyramids can even self-intersect. Here we show the worst possible representation of a saddle-shaped surface by a vertex star which itself is not even saddle-shaped.

$\mathcal{E}' = \sum \ell_e |\alpha_e|$. Both are a kind of bending energy. \mathcal{E}' has been used for the purpose of optimizing triangulations of surfaces (see [Alboul and van Damme 1996] and follow-up papers), and also of volumes [Dyn et al. 2001]. Its usage for formfinding was proposed by [Garanzha 2010].

Our work also involves fairness functionals operating on surfaces. These are frequently defined in terms of curvatures. We employ the *total absolute curvature* defined as the surface integral $\int |\kappa_1| + |\kappa_2|$, where κ_1, κ_2 refers to principal curvatures. Several deep theorems concern the gradient flow of such functionals. Well-studied examples include the surface area functional, and also the total mean curvature $\int \frac{\kappa_1 + \kappa_2}{2}$. Moving the surface with speed proportional to mean curvature $\frac{\kappa_1 + \kappa_2}{2}$ resp. Gauss curvature $\kappa_1 \kappa_2$ corresponds to L^2 -steepest descent of surface area resp. total mean curvature. It is known that in both cases, convex surfaces flow to round points [Andrews 1999; Huisken 1984]. Flows used for geometry processing purposes include the gradient flow of Willmore energy $\int (\kappa_1 - \kappa_2)^2$, see [Bobenko and Schröder 2005] for a mesh version. A sign-corrected Gauss curvature flow has been used by [Zhao and Xu 2006].

In the present paper, we do not use the gradient flow for minimizing total absolute curvature, but a method very similar to the one of Pellis et al. [2017].

A mesh energy may or may not be a discretization of a surface fairness functional. It is well known that discrete mean curvature $\frac{1}{2} \sum \ell_e \alpha_e$ is a discretization of total mean curvature $\int \frac{\kappa_1 + \kappa_2}{2}$, see [Cohen-Steiner and Morvan 2003]. It must be emphasized that our energies $\mathcal{E}, \mathcal{E}'$ do not enjoy such a property for general meshes, despite being called total “absolute mean curvature” in the literature.

Returning to curvature-based functionals, both our total absolute curvature $\int |\kappa_1| + |\kappa_2|$ and the functional $\int (\kappa_1^2 + \kappa_2^2)^{1/2}$ have an interpretation as *total variation* of the normal vector field. The latter functional is studied in detail in a recent preprint by Bergmann et al. [2019]. They also discuss the energies $\mathcal{E}, \mathcal{E}'$ and touch upon several of our topics. Energy-minimal meshes (considered in this paper) are mentioned as an unsolved problem.

2 REPRESENTATION OF SADDLES IN MESHES

Experience shows that visible deficiencies in fairness of polyhedral meshes occur mainly in negatively curved regions, i.e., parts of the mesh where the underlying smooth reference surface is locally saddle-shaped. The visual appearance of the mesh depends heavily




Fig. 3. *Saddles and numerical differential geometry.* We display the “cotan” mean curvature field \tilde{H} for a mesh inscribed in the surface $z = x^2 - y^2/2$. The x, y coordinates of vertices are taken from a regular lattice, slightly perturbed. Vertices and face normals approximate the smooth reference geometry to roughly the same extent in both (a) and (b), but \tilde{H} is much better behaved in (a), where vertex stars are saddle-shaped. In (b), vertex stars are no saddles, like in Fig. 2c. Compared to (a), a much smaller deviation from a precisely regular sampling makes \tilde{H} break down.

on the normal vectors of faces – Fig. 1 shows how a reflective surface reveals the normal vectors’ high variation. For this reason we have a closer look at the immediate neighbourhood of a vertex and visualize the normal vectors. Fig. 2 illustrates a fact which we feel has not been sufficiently appreciated so far: there is a big difference between the local shape of a smooth surface and the local shape of a mesh. In particular the normal vectors of faces in the immediate neighbourhood of a vertex behave in a much more irregular way than they do in the smooth surface case.

2.1 Saddle-shaped meshing of saddle-shaped surfaces

Even if a mesh is a regular sampling of a smooth saddle-shaped surface, vertex stars are not guaranteed to have proper saddle form.

Here saddle-shapedness means that there is a test plane (which can be considered a tangent plane) through the central vertex, intersecting the vertex star in precisely 4 line segments. We require that these four segments are not contained in a common half-plane, and that the vertex star projects onto the tangent plane in a 1-1 manner.

Examples of proper saddles are shown by Figures 2b and 4, while the vertex star of Fig. 2c is neither convex nor a saddle.

Figure 3 demonstrates how evaluation of discrete differential quantities can go wrong even for a mesh precisely inscribed in a smooth surface, if the local geometry of the mesh does not reflect the local geometry of the surface. This experiment supports the claim that visually smooth meshes are better behaved with respect to the numerical differential geometry, even if the mesh and normals of faces approximate a smooth reference geometry to the a similar extent.

2.2 The normal pyramid

Consider a vertex v and neighbours w_1, w_2 etc., in that order – see Fig. 4. The normal vectors of faces vw_iw_{i+1} form the *normal pyramid*. Such normal pyramids have been studied by [Günther et al. 2017; Jiang et al. 2016] who were in particular interested in cases where it has no self-intersections. This is a criterion for smoothness of the mesh. We attach a local xyz coordinate system with origin in v , such that the tangent plane is the xy plane.

The edges vw_j are partitioned into four sets of edges lying below resp. above the tangent plane. The number of elements of these sets




Fig. 4. (a) Good saddle, where vw_1w_2 and vw_4w_5 are no inflection faces. (b), (c): Bad saddles where the number of edges above and below the tangent plane is 2, 2, 1, 1 resp. 3, 1, 1, 1. Besides, there are inflection faces not intersecting the tangent plane.

can be 1, 2, 1, 2 (Fig. 4a) or 2, 2, 1, 1 (Fig. 4b) or 3, 1, 1, 1 (Fig. 4c). We want to give a criterion which is easily checked and which ensures that the normal pyramid is free of self-intersections. We say a face vw_iw_{i+1} is an *inflection face*, see [Günther et al. 2017], if neighbours $vw_{i-1}w_i$ and $vw_{i+1}w_{i+2}$ lie on different sides of the plane which carries the face f . In the simplest case of valence 4, every face of a saddle is an inflection face. We now say that a valence 6 saddle is a *good saddle* if the numbers of edges above and below the tangent plane is 1, 2, 1, 2 (Fig. 4a) and the two faces not intersecting the tangent plane are no inflection faces.

PROP. 1. *In a “good” saddle point, the pyramid of vertex normals is free of self-intersections.*


PROOF. We cut off the normal pyramid by the plane η which has equation $z = 1$. Vertices w_1, \dots, w_6 are neighbours of v (indices modulo 6). In the plane η we mark the points $n_{i,i+1}$, which are defined as intersection with η of the normal of the plane $[vw_iw_{i+1}]$. Consider a ray $L(t)$ rotating around v but always contained inside the surface of the mesh. Here t is a time parameter. Consider the plane $L(t)^\perp$ orthogonal to $L(t)$ and trace the intersection $L'(t) = L(t)^\perp \cap \eta$.

There are six time instances t_1^*, \dots, t_6^* where $L(t_j^*)$ coincides with an edge vw_j . Then $L'(t_j^*)$ equals the line $n_{j-1,j} \vee n_{j,j+1}$ (yellow in the inset figure). In the time interval where $L(t)$ is contained in the face vw_jw_{j+1} , $L'(t)$ rotates about $n_{j,j+1}$. The rotation is consistently in one direction. Four times the line $L(t)$ lies in the tangent plane. We assume the vertices are numbered so that these four lines lie in faces vw_jw_{j+1} for $(j, j+1) = (2, 3), (3, 4), (5, 6), (6, 1)$ (which are inflection faces). This happens at time instances $t = t_{j,j+1}$. The corresponding line $L'(t_{j,j+1})$ passes through the point o where the z axis intersects η (shown in red).

Observe the ray $L(t)$ rotating from edge vw_4 to edge vw_5 . The edge vw_5 is reached at time $t = t_5^*$, before the plane $[vw_5w_6]$ is crossed (this is the non-inflection property of vw_4w_5). The line $L'(t)$ during this rotation has not crossed n_{56} . Similarly, during the rotation in the time interval $[t_5^*, t_6^*]$, the line $L'(t)$ does not cross n_{34} . Therefore n_{45} lies in the interior of the triangle o, n_{34}, n_{56} .

Likewise, n_{12} lies in the interior of the triangle o, n_{61}, n_{23} . \square

Remark. If a triangle mesh of regular combinatorics is generated by sampling a smooth surface in a geometrically regular way, it is hard to create saddle points with more than 4 inflection faces, and






Fig. 5. *Design of triangle meshes.* (a) We interactively design meshes with the principal-symmetric property. (b) Their local symmetries respect curvatures, so their normal pyramids resp. Gauss image has symmetric shape and is free of self-intersections. (c) The principal-symmetric property means in particular that one third of edges is aligned with principal curvatures, and vertex normals form developable strips.

typically also the asymmetric situations of Fig. 4b,c will not occur. There is however a nonzero probability that the mesh will have vertices which are neither convex nor proper saddles like in Fig. 2.

2.3 Principal-symmetric triangle meshes

Arguably the best way to resolve a saddle-shaped reference surface Φ as a mesh is that in each vertex, the local geometry of the mesh around a vertex v has the same symmetries as Φ . For a typical valence 6 vertex v this means that two edges through v , say vw_1 , vw_4 (see Figure 6), are aligned with a principal curvature direction, while the others are symmetric to it. Differential geometry tells us that they are then aligned with two curves enjoying the same normal curvature. To discretize these two properties we must distinguish between principal edges and non-principal edges and require that in each valence 6 vertex, there are exactly two principal edges which in the vertex star are opposite to each other. We impose the following set of conditions.




Fig. 6. *The spherical condition imposed on principal-symmetric meshes.* Vertex v is endowed with a sphere $\Sigma(v)$ containing v itself and 4 neighbours w_2, w_3, w_5, w_6 . In case of negative curvature we require that the remaining vertices lie on the other side of the sphere's tangent plane.

- A valence 6 vertex has four non-principal edges vw_j . According to [Schling et al. 2018], equality of normal curvatures is discretized by the property that v and all involved w_j 's lie on a sphere $\Sigma(v)$. See Fig. 6.
- Every vertex v is now endowed with a normal vector $n(v)$, namely the normal vector of the sphere $\Sigma(v)$. For a principal edge vw_j we require that $n(v), n(w_j), w_j - v$ are co-planar. In this way we construct developable strips orthogonal to the mesh along all the principal edges, see Fig. 5.
- In a valence 6 vertex v with angle sum greater than 2π we assume a saddle shape. Then the sphere $\Sigma(v)$ and a principal edge vw_j are required to lie on different sides of the sphere's tangent plane.
- Vertices of valence different from 6 are considered singularities, comparable to umbilic points, and no conditions are imposed there.

These *principal-symmetric* conditions are rather strong, but they yield the best results regarding shape reproduction by meshes (for the meaning of 'best', see the start of this paragraph). We use them for interactive design of triangle meshes, see Figures 5 and 14. For implementation see §5.

3 ENERGIES OF POLYHEDRAL SURFACES

The local shape properties of meshes discussed in §2 are in direct relation to the *variation of normal vectors* of meshes. E.g. the total variation of normal vectors around a vertex, which equals the sum of unsigned dihedral angles of the vertex star, can be seen as a kind of spherical perimeter of the normal pyramid. A smaller total variation corresponds to a smaller perimeter which is obviously achieved in cases where the normal pyramid is free of self-intersections. We therefore propose to construct a discrete mesh energy \mathcal{E} as explained below, which penalizes big dihedral angles and in turn, a higher total variation of normals.

Consider a polyhedral surface $P = (V, E, F)$ with vertex set V , edge set E and face set F . Each face $f \in F$ has a normal vector \mathbf{n}_f . Any edge $e \in E$ has a length ℓ_e , a dihedral angle α_e , and a dual edge length ℓ_e^* , which is defined by the relations

$$e = f \cap f' \implies |\alpha_e| = \angle(\mathbf{n}_f, \mathbf{n}_{f'}), \ell_e^* = \|\mathbf{n}_f - \mathbf{n}_{f'}\| = 2 \sin \frac{\alpha_e}{2}.$$

The energy

$$\mathcal{E}'(P) = \sum_{e \in E} \ell_e |\alpha_e|$$

is built similar to discrete mean curvature and penalizes big dihedral angles (i.e., a big contribution to the total variation of normal vectors). A slight modification leads to the energy

$$\mathcal{E}(P) = \sum_{e \in E} \ell_e \ell_e^* = 2 \sum_{e \in E} \ell_e \left| \sin \frac{\alpha_e}{2} \right|,$$

which is easier to handle from the computational viewpoint.

With $\sin x \approx x$ for small values of x , the energies $\mathcal{E}, \mathcal{E}'$ assume similar values; they act identically for all practical purposes.

3.1 Differential-geometric interpretation of the energy \mathcal{E}

Recall the definition of the shape operator of a surface Φ [do Carmo 1976]. When sitting in a point $x \in \Phi$ and moving by the small amount δx , the unit normal vector n is incremented by $\delta n \approx -W \delta x$. The shape operator W is linear. Its eigenvectors are orthogonal and indicate the principal directions: If the increment δx is in the j -th *principal* direction ($j = 1, 2$), then the normal vector increment is

$$\delta n = -\kappa_j \delta x,$$

where κ_j is the j -th principal curvature. In an orthonormal coordinate system aligned with the principal directions, W is described by the diagonal matrix $(\kappa_1 \kappa_2)$.

For a *principal quad mesh* $P = (V, E, F)$, the edge polylines approximate the principal curvature lines of a smooth surface Φ . In such cases the faces of the mesh (away from umbilics) are near-flat, and can in fact be assumed flat [Bobenko and Suris 2009].




Fig. 7. Convergence of mesh energy to total absolute curvature $C(\Phi)$. (a)–(c) show finer and finer principal meshes, each approximating the same reference shape Φ . The energies approach $C(\Phi)$. Every edge e is color-coded according to $\ell_e |\alpha_e|$. The reference surface Φ is represented by the triangle mesh (d).

For any face f of the mesh, the normal vector \mathbf{n}_f is also a normal vector of the underlying smooth reference surface Φ in an appropriate point. This location, when projected back onto the face f , yields a point c_f , see Figure 8.

For any edge $e = f \cap f'$, we now have the following correspondences between the discrete mesh situation and the smooth surface situation:

$$\begin{aligned} c_{f'} - c_f &\approx \delta x, \quad \mathbf{n}_{f'} - \mathbf{n}_f \approx \delta n, \quad \ell_e^* \approx \|\delta n\| \approx |\kappa_e| \|\delta x\| \\ \Rightarrow \mathcal{E}(P) &= \sum_e \ell_e \ell_e^* \approx \sum_e |\kappa_e| \ell_e \|c_f - c_{f'}\|. \end{aligned}$$

Here κ_e is the principal curvature in the direction orthogonal to e . The quadrangle spanned by $c_f, c_{f'}$ and e has area approximately $\frac{1}{2} \|c_f - c_{f'}\| \cdot \ell_e$. One half of the edges corresponds to the first principal direction, the other half to the 2nd principal direction, so the term $\ell_e \|c_f - c_{f'}\|$ occurring in the formula above is the area of influence of the edge e . Summing up,

$$\mathcal{E}(P) \approx \int_{\Phi} (|\kappa_1| + |\kappa_2|) dA = C(\Phi).$$

Numerical experiments confirm this, see Fig. 7. We return to the functional $C(\Phi)$ in §4.

3.2 Energy gradient

We now investigate energy-minimal meshes, and in particular meshes which have minimal energy subject to the constraint that they approximate a given smooth surface Φ . We will eventually arrive at the conclusion that principal meshes are energy minimizers.

Triangle meshes with fixed connectivity form a linear space identifiable with $(\mathbb{R}^3)^{|V|}$. The energy \mathcal{E} as a function of vertices $v \in V$ has gradient $(\nabla \mathcal{E}_v)_{v \in V}$ if for a smooth path $(v(t))_{v \in V}$ the energy changes according to $\frac{d}{dt} \mathcal{E} = \sum_{v \in V} \nabla \mathcal{E}_v^T \frac{dv}{dt}$. Certainly in energy-minimal meshes, all individual vectors $\nabla \mathcal{E}_v$ of free vertices v must vanish. If we consider the class of meshes whose vertices are constrained to a smooth surface Φ , energy-minimal meshes have the property that $\nabla \mathcal{E}_v$ is orthogonal to Φ .

A difficulty arises from the fact that \mathcal{E} is not differentiable for zero edge length and zero dihedral angles. The latter case cannot be excluded and in fact typically is present in optimal meshes. We start our discussion by describing the energy gradient in the generic case. The proof of the following result is contained in the appendix.




Fig. 8. To interpret the energy \mathcal{E} , we look at faces f, f' intersecting in an edge e . Stepping from f to f' (actually, from point c_f to point $c_{f'}$), we increment position by $\delta x = c_{f'} - c_f$ and the normal vector by $\delta n = \mathbf{n}_{f'} - \mathbf{n}_f$.

PROP. 2. For a triangle mesh $P = (V, E, F)$ with edge lengths and dihedral angles nonzero, the energy \mathcal{E} is differentiable, with

$$\begin{aligned} \nabla \mathcal{E}_v &= \sum_{e=vw} \frac{\ell_e^*}{\ell_e} (v - w) + \\ &+ \sum_{\substack{e=vw^+ = f \cap f' \\ f = vw^+ w^-, f' = vw^- w^-}} \ell_e \left(\frac{w^+ - w}{2|f|} \times \frac{\pi_f(\mathbf{n}_{f'})}{\ell_e^*} + \frac{w - w^-}{2|f'|} \times \frac{\pi_{f'}(\mathbf{n}_f)}{\ell_e^*} \right) \\ &+ \sum_{f = vw^+ w^+, \bar{e} = ww^+ = f \cap \bar{f}} \ell_{\bar{e}} \frac{w^+ - w}{2|f|} \times \frac{\pi_f(\mathbf{n}_{\bar{f}})}{\ell_{\bar{e}}^*} \end{aligned} \quad (1)$$

(for notation see Fig. 9). Each sum is over the immediate neighbours w of the central vertex v , and involves the area $|f|$ of faces and the orthogonal projection π_f onto the 2D subspace defined by the face f .

Ideal energy-minimal meshes I. We return to the case of meshes approximating a reference surface Φ . It is easy to identify a case where $\nabla \mathcal{E}_v$ is orthogonal to Φ : This surely happens if the neighbourhood of v which is involved in computing $\nabla \mathcal{E}_v$ is symmetric w.r.t. a half-turn about the surface normal in v . In that case also $\nabla \mathcal{E}_v$ has this symmetry, i.e., is orthogonal to Φ . The vertex v is then in an energy-optimal position on the surface Φ insofar as $\nabla \mathcal{E}_v$ has no component tangential to Φ .

Of course this ideal symmetric situation cannot be expected to happen in the strict sense, but any triangle mesh which approximates a smooth net of curves on a smooth surface is close to it.

One-sided energy gradient and ideal minimal case II. In the case of vanishing dihedral angles, the energy still has a one-sided gradient, and it is useful to identify situations where the one-sided energy gradient is *positive*. This is guaranteed if each individual contribution to the derivative of energy is positive, i.e., if the mesh is deforming with time $t, t \geq 0$, we have

$$\lim_{t \searrow 0} \frac{1}{t} \left(\ell_e(t) \ell_e^*(t) - \ell_e(0) \ell_e^*(0) \right) > 0. \quad (2)$$

We can show this is true in case $e = f \cap f'$ where $f \cup f'$ is a parallelogram. The proof of this statement is contained in the appendix. Such a situation is highly idealized, but any polyhedral mesh which approximates the parameter lines on a smooth surface consists of quads which are approximately parallelograms, see e.g. Fig. 10.




Fig. 9. Notation for Prop. 2. Edges e emanate from the central vertex v , edges \bar{e} lie opposite to it. Faces f, f' lie in the 1-ring neighbourhood, faces \bar{f}, \bar{f}' are adjacent to it.




Fig. 10. *Energy minimization: Ideal situation II.* Here a surface is endowed with a principal quad mesh, each quad having approximately rectangular shape. Thus we are close to ideal situation II encountered in our discussion of energy-minimizing meshes (in fact this meshing is energy-minimal). In addition to the meshing, the reference surface itself has minimal total absolute curvature, for prescribed boundaries. The rendering is that of a freeform glass facade.

3.3 Energy-minimal meshes

Previous ideal cases lead to the following conclusion:

ENERGY-MINIMAL MESHES. *Among meshes which approximate a given reference surface Φ such that edges approximate a smooth network of curves on Φ , the principal meshes enjoy minimal energy. In negatively curved areas, other meshes have higher energy. The energy of a minimizer equals the total absolute curvature of Φ .*

We argue as follows. Consider a mesh (V, E_0, F) with flat quadrilateral faces and make it a triangle mesh $(V, E_0 \cup E_1, F')$ by introducing diagonals in all faces (the set of diagonals is E_1). By the smooth curve network assumption we are close to ideal situations I and II discussed above. The energy is split into contributions $\mathcal{E} = \mathcal{E}_0 + \mathcal{E}_1$ from edge sets E_0, E_1 , respectively. From ideal situation I we know that in any vertex, $\nabla \mathcal{E}_0, v$ is orthogonal to Φ . Thus \mathcal{E}_0 is infinitesimally constant for all infinitesimal changes of vertices tangential to Φ . From ideal situation II we know that changing the mesh away from planarity of faces increases \mathcal{E}_1 .

We now apply the following argument. We replace Φ by a local 2nd order Taylor approximation, namely a paraboloid with equation $z = \alpha x^2 + \beta y^2$. We then show that among all parallelogram vertex stars of the origin which are inscribed in the paraboloid, the value $\sum \ell \ell^*$ per area is minimal for rectangles. We skip this lengthy but




Fig. 11. Finding an energy minimum. In our argumentation we check “energy per area” for meshes inscribed in paraboloids $z = \alpha x^2 + \beta y^2$ whose faces are parallelograms. At right we show this value as a function of the polar angle of an edge. We get the result that the energy is minimal if the polar angle is a multiple of $\pi/2$, where faces are actually rectangles.




Fig. 12. *Interactive design of small-energy meshes.* An initial shape (left) is modified interactively by pulling at handles. A fairness functional together with energy \mathcal{E} guides the shape of the final mesh (right). We also show the Gauss images of meshes (unit normal vectors of faces). Self-intersections are avoided, which would not be the case if we neglected \mathcal{E} .

elementary computation — see Fig. 11 for a demo. Only in case $\alpha\beta < 0$ (negative curvature) the minimum is significant enough, however. This is in accordance with the fact that if $\alpha\beta > 0$, \mathcal{E} discretizes total mean curvature, independent of the meshing.

Summing up, minimal energy is achieved for polyhedral meshes whose faces are approximately rectangular, i.e., which are principal quad meshes. In case of negative curvature, other meshes have higher energy. The chain of arguments presented above is not proof, but numerical evidence so far supports the statement.

The problem of approximating a given smooth surface by an energy-minimal mesh is highly nonlinear and we cannot expect to solve it without knowing in advance the combinatorics of the solution. However it makes still sense to improve the energy \mathcal{E} by continuous optimization, and also to include it in an interactive design tool. This is because minimizing the energy \mathcal{E} modifies a mesh such that the individual normal pyramids typically have no self-intersections (see Fig. 12). This experimental result is entirely plausible since the energy \mathcal{E} penalizes a larger spherical perimeter of normal pyramids.

4 TOTAL ABSOLUTE CURVATURE OF SURFACES

We already saw that the total absolute curvature

$$C(\Phi) = \int_{\Phi} (|\kappa_1| + |\kappa_2|) dA$$

of a smooth surface Φ coincides with the smallest achievable value of the discrete energy \mathcal{E} for polyhedral meshes approximating Φ . It therefore makes sense to ask the question which surface Φ , under given boundary conditions, achieves the lowest value of $C(\Phi)$.

The shape operator W in a principal coordinate frame is described by the diagonal matrix $(\kappa_1 \kappa_2)$, so total absolute curvature equals

$$C(\Phi) = \int \|W\|_1, \quad \text{where} \quad \|W\|_1 = |\kappa_1| + |\kappa_2|$$

is the 1-norm (trace norm) of the shape operator. It can be used to bound the infinitesimal normal vector increment $\delta n \approx -W\delta x$ caused by the position increment δx (cf. the discussion in the beginning of §3). Like any proper matrix norm, $\|W\|_1$ obeys

$$\|\delta n\| \approx \|W\delta x\| \leq \|W\|_1 \cdot \|\delta x\|. \quad (3)$$

Because of this formula, the total absolute curvature $C(\Phi)$ can be interpreted as the total variation $TV(n)$ of the normal vector field, if the size of derivatives is measured using $\|\cdot\|_1$.

Besides $\|W\|_1$, also other functions of the principal curvatures obey inequality (3) and yield a sensible curvature measure. Bergmann et al. [2019] define total variation as $\int \|W\|_2$, where $\|W\|_2 = (\kappa_1^2 + \kappa_2^2)^{1/2}$ is the Frobenius norm of the shape operator. We use $\|W\|_1$ because of its relation to the discrete mesh energy \mathcal{E} .

Only in special situations we are able to describe the minimizers of $C(\Phi)$. Consider e.g. a surface with rotational symmetry about the z axis described by a radius function $r(z) > 0$ and boundary values $r(z_0) = r_0, r(z_1) = r_1$. It is not difficult to compute the area element $dA = 2\pi r \sqrt{1 + r'^2} dz$ and curvatures $\kappa_1 = 1/(r\sqrt{1 + r'^2})$, $\kappa_2 = r''(1 + r'^2)^{-3/2}$. Thus, $\int |\kappa_1| dA = \int 2\pi dz = 2\pi(z_1 - z_0)$ depends only on the boundary conditions, whereas $\int |\kappa_2| dA = \int |r''|(\dots) dz$ vanishes if and only if $r(z)$ is linear. Thus, among rotational surfaces, minimizers are cones and cylinders (see Fig. 13).

The computation of surfaces which minimize energy under different boundary conditions is treated below. Examples are shown throughout this paper (Figures 1, 10, 13, 15, 16, 21).

Relation between curvature functionals and mesh energies. It is known that for a mesh P approximating a smooth surface Φ , under reasonable assumptions the expression $\sum_e \ell_e \alpha_e$ approximates $\int_\Phi \kappa_1 + \kappa_2$, see [Cohen-Steiner and Morvan 2003]. As a consequence, $\sum_e \ell_e \alpha_e$ is approximately independent of the meshing and only dependent on the geometric shape described by the mesh. No such statement is true for the energies $\mathcal{E}, \mathcal{E}'$ considered by this paper, even if the definition $\mathcal{E}'(P) = \sum \ell_e |\alpha_e|$ is very similar to the above. $\mathcal{E}(P)$ changes with remeshing, and only if P is principal, the value $\mathcal{E}(P)$ can be identified with the functional $\int |\kappa_1| + |\kappa_2|$, see Fig. 7.

One must beware of the following argument, which is correct as far as it goes, but can easily lead to a wrong conclusion. Rounding off a mesh P with cylinders along edges and small suitable surface pieces near vertices will create a smooth surface Ψ with $C(\Psi) \approx \mathcal{E}(P)$. This does *not* imply that in general, for meshes P which approximate a smooth surface Φ , we have $\mathcal{E}(P) \approx C(\Phi)$. A corresponding statement on the energy \mathcal{E}' and the functional $C_2(\Phi) = \int_\Phi (\kappa_1^2 + \kappa_2^2)^{1/2}$ is made by Bergmann et al. [2019]. It does not imply that \mathcal{E}' discretizes C_2 in general.




Fig. 13. Surfaces minimizing total absolute curvature. Here we show surfaces with prescribed boundary curves which minimize $\int |\kappa_1| + |\kappa_2|$. Among rotational surfaces, minimizers are cylinders and cones. Observe that minimizers have features not seen in minimizers of e.g. Willmore energy.

5 IMPLEMENTATION

We are solving three different computational problems:

- Finding surfaces of small total absolute curvature under given boundary conditions. Such surfaces are represented by meshes, which have to be optimized w.r.t. a nonlinear target functional.
- Optimizing meshes towards smaller energy. This can be done such that the given mesh is required to approximate an unchanging reference surface. In any case this procedure means moving the vertices in order to improve the value of the energy, without changing the connectivity (which usually prevents us from reaching energy values equal to the total absolute curvature).
- Interactive design with principal-symmetric meshes.

All three tasks boil down to solving a potentially large system of nonlinear constraints, and finding that point in the solution variety where a certain target functional, as well as standard fairness functionals, assume small values. For that purpose, we employ *guided projection* as proposed by Tang et al. [2014]. The procedure works well for sparse systems of polynomial equations of degree at most 2. This low degree is mostly achieved by introducing auxiliary variables or computing with approximations. The next paragraphs describe our setup in greater detail.

5.1 Computing energy-minimal meshes

We discuss how to optimize meshes such that the energy $\mathcal{E} = \sum_{e \in E} \ell_e \ell_e^*$ is minimal and the vertices of the mesh obey constraints like proximity to a reference surface, or proximity to handles moved by the designer, see Fig. 12. We confine ourselves to optimizing vertex positions without changing combinatorics.

Setup of variables and constraints. For a mesh (V, E, F) , the collection “ X ” of variables consists of the coordinates of vertices, edge lengths ℓ_e , dual edge lengths ℓ_e^* , and normal vectors \mathbf{n}_f of faces. The constraints “ $F(X) = 0$ ” imposed on these variables are the quadratic defining relations for ℓ_e and ℓ_e^* together with $\ell_e = \delta_e^2$, $\ell_e^* = \mu_e^2$ to ensure that ℓ_e, ℓ_e^* are nonnegative. The normal vectors of faces are defined by their orthogonality to edges ($\mathbf{n}_f^T(v - w) = 0$ whenever $v, w \in f$) and by the normalization $\|\mathbf{n}_f\|^2 = 1$. We do not attempt to eliminate redundant constraints since the method is resilient w.r.t. redundancies and they even appear to stabilize the solution.

In addition to these constraints, we have the condition that vertices are confined to a smooth reference surface Φ . Such soft constraints are linearized by replacing Φ by its tangent planes, appropriately updated as iteration progresses, cf. [Tang et al. 2014]. Boundary curves are treated in a similar manner.

The guided projection method, in each round of iteration, computes a linearization “ $AX = B$ ” of the constraints and solves $\|AX - B\|^2 + R(X) \rightarrow \min$, where $R(X)$ is a quadratic regularizer,

$$R(X) = w_{\text{reg}} \|X - X^{\text{prev}}\|^2 + w_{\text{fair}} E_{\text{fair}}(X) + w_{\text{target}} E_{\text{target}}(X).$$

Here X^{prev} is the value of X in the previous round of iteration. E_{fair} is a fairness functional. We use a sum of squares of second differences of mesh polylines. The third contribution to R , the target functional, here is $E_{\text{target}}(X) = \mathcal{E} = \sum \ell_e \ell_e^*$. Its weight has to be chosen such that R is still positive definite. When employed in a design situation,




Fig. 14. *Mesher with good reflections.* Starting from a triangle mesh representing a torus (a), we interactively design a mesh with the principal-symmetric property (b). It features normal developables along the principal edges (c). Subdivision and subsequent optimization produces finer principal-symmetric meshes. We show how reflections in the finer triangle meshes (d,e) compare with reflection in a smooth reference surface Φ (f) and a principal quad mesh approximating Φ (g). These results confirm that both principal quad meshes and principal-symmetric triangle meshes provide a good discrete representation of local surface geometry, and reflection patterns of the smooth reference geometry are faithfully reproduced.

Subfigure (h) shows another design. A mesh (visible only as edges) is modified such that the highlighted vertices are fixed in order to create a nodoidal-shaped mesh. The principal-symmetric condition is enforced only in the highlighted polylines, otherwise the energy \mathcal{E} is minimized. The result is subdivided and optimized again. We show reflection in the coarse version and a smooth reference shape. Faithful reflections are much less evident than for (d) and (e).

the designer's input is added in the form of appropriately weighted constraints.

5.2 Computing principal-symmetric meshes

Like for energy-minimization, implementation of a full-fledged design tool for meshes enjoying the principal-symmetric property is not straightforward. This is because such meshes must follow the principal curvature lines of the mesh's geometric reference shape Φ , which is extremely sensitive w.r.t. shape changes. A design tool therefore must handle changes in combinatorics. We did not attempt to implement design in such generality. Therefore the procedure described in this section is only part of the overall design pipeline.

Setup of variables and constraints. This follows directly from the description of the principal-symmetric property in §2. Variables are vertices v , sphere centers $c(v)$, and auxiliary unit normal vectors $\mathbf{n}(e)$ for planes associated with the principal edges. For any co-spherical neighbour w of v we need the constraint $\|v - c(v)\|^2 = \|w - c(v)\|^2$, and for any principal edge vw we need the coplanarity constraints $\mathbf{n}(e)^T(v - w) = \mathbf{n}(e)^T(c(v) - v) = \mathbf{n}(e)^T(c(w) - w) = 0$. We also require that $\mathbf{n}(e)^T \mathbf{n}(e) = 1$. The saddle-shaped constraint requires that for any vertex v with principal neighbour w , we have

$$(v - c(v))^T(v - w) < 0$$

whenever that vertex has negative curvature. We therefore require that $(v - c(v))^T(v - w) = -\delta_{v,w}^2$, where $\delta_{v,w}$ is a dummy variable associated with the edge vw . The linearization of proximity constraints has already been mentioned above.




Fig. 15. Minimizers of total absolute curvature where boundaries plus tangent planes along boundaries are prescribed.

This concludes the setup of variables and constraints. Regularizing terms are the same as in the previous paragraph, with w_{target} nonzero if in addition to the principal-symmetric property we also want to achieve small energy. For results see Figure 14.

5.3 Minimizing total absolute curvature

We show how to compute a surface Φ which minimizes total absolute curvature under boundary conditions. The surface is represented by a mesh (V, E, F) of sufficiently high resolution.

Before we list variables and constraints, we describe how we discretize the required differential quantities of Φ . To estimate the shape operator, we use the *normal cycle* method of Cohen-Steiner and Morvan [2003]. They extend the shape operator to a symmetric 3×3 matrix by requiring that it kills the normal vector of the surface. Their formula for the shape operator W_v , aggregated over the star of a vertex v , involves the dihedral angles α_e of edges emanating from v . By replacing α_e with $\sin \alpha_e$ we get a modified version W'_v of this shape operator:

$$W'_v = \frac{1}{2} \sum_{e=vw=f \cap f'} (\mathbf{n}_f \times \mathbf{n}_{f'}) (w - v)^T \in \mathbb{R}_{\text{symm}}^{3 \times 3}.$$

In this way, the total absolute curvature is discretized as

$$C(\Phi) \approx \sum_v |\lambda_{1,v}| + |\lambda_{2,v}|,$$

where the numbers $\lambda_{j,v}$ indicate the eigenvalues of W'_v corresponding to the two tangential eigenvectors. Their geometric meaning is principal curvatures times the area of influence of the vertex v .

Setup of variables and constraints. We use as variables the coordinates of vertices v , the face normal vectors \mathbf{n}_f , and their cross products $\mathbf{n}_f \times \mathbf{n}_{f'}$ for all faces f, f' which share an edge. Further, we have matrices W'_v with their eigenvalues $\lambda_{j,v}$ and eigenvectors $\mathbf{d}_{j,v}$ ($j = 1, 2$). These variables are subject to constraints:

- (1) If $v, w \in f$, then $(v - w)^T \mathbf{n}_f = 0$. We ensure that all \mathbf{n}_f point to the same side of the mesh.
- (2) $\|\mathbf{n}_f\|^2 = \|\mathbf{d}_{1,v}\|^2 = \|\mathbf{d}_{2,v}\|^2 = 1$.

- (3) For each edge $e = vw = f \cap f'$ we define $s(v, w) = \mathbf{n}_f \times \mathbf{n}_{f'}$ (if the edge is traversed from v to w , then f' is to the right, and f is to the left).
- (4) $W'_v = \frac{1}{2} \sum_{e=vw} s(v, w)(w - v)^T$, by definition.
- (5) $W'_v \mathbf{d}_{j,v} = \lambda_{j,v} \mathbf{d}_{j,v}$, for all vertices $v \in V$, and $j = 1, 2$.
- (6) To ensure that eigenvectors $\mathbf{d}_{j,v}$ are tangential to the mesh, we require $\mathbf{d}_{j,v}^T \sum_{f \ni v} \mathbf{n}_f = 0$ for all vertices $v \in V$.
- (7) Eigenvectors are orthogonal: $\mathbf{d}_{1,v}^T \mathbf{d}_{2,v} = 0$ for all $v \in V$.
- (8) The third eigenvalue of W'_v vanishes; we therefore require $\text{tr}(W'_v) = \lambda_{1,v} + \lambda_{2,v}$, for all $v \in V$.

Here (3), (4) are explicit definitions. Items (1)+(2) implicitly define normal vectors, and (2)+(5) do the same for eigenvectors and eigenvalues of W'_v . The remaining equations can be considered constraints. Further constraints are given by boundary conditions like fixed vertices or vertices gliding on boundary curves. They are linearized as described above. We use the same regularizer $w_{\text{reg}} \|X - X^{\text{prev}}\|^2$ as in previous paragraphs, and as an additional fairness functional, a sum of squares of umbrella vectors [Kobbelt et al. 1998].

For target functional $C(\Phi)$ we use a weight iteration approach. In the i -th round of iteration we do not wish to minimize $\sum |\lambda_{1,v}| + |\lambda_{2,v}|$ directly, but rather the quadratic function

$$E_{\text{target}} = \sum_{v \in V} \frac{\lambda_{1,v}^2}{|\lambda_{1,v}^{(i-1)}| + \epsilon} + \frac{\lambda_{2,v}^2}{|\lambda_{2,v}^{(i-1)}| + \epsilon}.$$

Here $\lambda_{j,v}^{(i-1)}$ is the value of the variable $\lambda_{j,v}$ assumed in the previous round of iteration. E_{target} is part of the regularizer $R(X)$.

6 RESULTS AND DISCUSSION

The results and procedures contained in this paper provide tools for both formfinding and design.

Formfinding via total absolute curvature. For formfinding, we have at our disposal the procedure to find a surface minimizing total absolute curvature for given boundary conditions. In addition to that, a mesh approximating that surface has minimal energy, if its faces are planar quads following the principal curvature lines. Meshing and postprocessing, however, are not contributions of this




Fig. 16. Two-step architectural designs. Firstly we find surfaces which minimize total absolute curvature, secondly we mesh these surfaces along principal curves, making the energy minimal.




Fig. 17. Pipelines for (i) computing surfaces with minimal total absolute curvature, (ii) computing meshes with small energy which approximate a fixed design surface, and (iii) meshing along principal curvature lines, which yields the energy-minimal meshes approximating the design surface.

paper. We refer to [Bommes et al. 2009] for meshing with edges aligned with principal curvatures, and to [Tang et al. 2014] for postprocessing for planarity of faces. Examples are given by Figures 1, 10, 13, 15, 16, 21.

Interactive design of meshes. Principal-symmetric meshes are those triangle meshes which exhibit the most faithful representation of a surface's local shape. The geometric modeling tool described by §5 is incorporated into a design pipeline as follows: (i) Design a desired geometric shape Φ . (ii) Compute the cross field of principal directions on Φ and create a quad mesh which follows this cross field, using e.g. the method of [Bommes et al. 2009]. Suitable diagonals in faces of this mesh now yield a triangle mesh whose combinatorics is already that of a principal-symmetric triangle mesh. (iii) Impose the principal-symmetric property and interactively modify the mesh further, always maintaining the principal-symmetric property.

We do not regard (i)+(ii) as new contributions. Task (iii) is solved by *guided projection* as described in §5. The unchangeable combinatorics of the principal curves of the mesh during phase (iii) can be seen as a restriction of design freedom, but on the other hand it is also a guarantee that essential shape characteristics cannot be destroyed by the designer. Results are shown by Fig. 14.

The design of meshes with a small energy \mathcal{E} (see Fig. 18) is much more straightforward: One uses the *guided projection* procedure directly, in the same manner as proposed by [Tang et al. 2014].

Implementation details. Figure 19 gives statistics for minimizing total absolute curvature, and Fig. 20 does the same for mesh design. Examples involving the principal-symmetric property (Figures 5, 14) have been constructed by a multiresolution approach: We first design a coarse mesh, where the response time of the algorithm




Fig. 18. *Mesh optimization.* The hybrid mesh with triangular and planar quadrilateral faces used for the roof of the Islamic Art exhibition in the Louvre, Paris, undergoes energy minimization while retaining proximity to the original. The final mesh is overlaid by the original in red. Quads are kept planar via constraints $\ell_e^* = 0$, for all diagonals e of planar quads.

Fig.	V	F	E	# var	w_{reg}	w_{fair}	w_{target}	w_{prox}	$C(\Phi)$	# it.	T_{total}
1	2784	5376	8160	87k	10^{-3}	.30	.05	—	94.3	18	58.3
10	2606	4992	7597	82k	10^{-3}	.20	.03	—	179.3	11	32.8
13a,16b	1056	2048	3104	33k	10^{-3}	.30	.05	—	124.4	14	16.1
13b	4224	8193	12416	133k	10^{-3}	.20	.10	—	180.2	11	57.4
13c	555	1017	1572	17k	10^{-3}	.20	.05	—	45.5	7	3.8
15left	6084	11949	18032	193k	10^{-3}	.05	.01	—	36.9	8	69.4
15right	2166	4236	6403	68k	10^{-3}	.10	.05	—	171.6	17	45.3
21	2118	3991	6112	66k	10^{-3}	.20	.05	—	143.8	17	39.2


(Code in Python, on Intel® Core™ i7-6700HQ CPU with 2.6 GHz, 16GB RAM)

Fig. 19. Statistics for minimizing $C(\Phi)$. We give the number of vertices etc. of the meshes we ran the optimization on (not necessarily the meshes shown in figures), as well as the weights used for regularizing functionals in the *guided projection* method. w_{prox} is the weight of soft constraints of vertices' proximity to a reference surface. We show total computation time in seconds. Experience shows that after about half the number of iterations shown, the minimizer's shape is more or less defined.

is small (the designer has to wait at least one iteration before continuing, and typically the shape stabilizes after 25 or so iterations). Subsequently we apply subdivision and optimize again. For the finer mesh, the response time is considerably longer, but only few iterations are needed.

Limitations. An obvious limitation of the design methods presented here is that the combinatorics is fixed during the process. This has already been discussed above. Another limitation is the computation time. Calling the design method 'interactive' surely is borderline (then again, certain prospective users are currently happy to wait hours for a genetic design to emerge).

Thirdly, we are concerned about the regularity of the surfaces and meshes which are the result of optimization. The features present in Fig. 16b or Fig. 21 suggest that it is advisable to include fairness functionals in our optimization. Indeed, the minimizers of total absolute curvature will exhibit cone points if a boundary point is kept fixed (because rounding off a cone's vertex with radius r yields absolute curvature of magnitude $O(1/r)$ concentrated on an area of magnitude $O(r^2)$, which is $O(r)$ in total, i.e., negligible).



Conclusion. We have presented a detailed study of the local geometry of vertex stars and its implications on visual smoothness of polyhedral surfaces. Besides theoretical results concerning the shape of the normal pyramids, we present certain nonlinear functionals which quantify the variation of the normal vector field over an entire surface (both polyhedral and smooth). We analyze these functionals and demonstrate optimality of principal meshes in the discrete case. We show how to utilize these functionals for formfinding and design at interactive or near-interactive speeds. Design further is performed on basis of the principal-symmetric property of meshes, which is a strong condition ensuring good discrete representation of saddle points.

Future Research. One direction of future research is the connection between triangle meshes with "nice" hexagonal normal pyramids on the one hand, and meshes with flat faces of "nice" hexagonal shape on the other hand. Work by Günther et al. [2017] about projective dualities and by Jiang et al. [2015] on polyhedral patterns

Fig.	V	F	E	# var.	w_{reg}	w_{fair}	w_{target}	w_{prox}	\mathcal{E}	T/iter
5ac	714	2058	1344	25k	10^{-6}	1.0	10^{-6}	—	16.0	0.53
12	189	525	336	6k	10^{-6}	1.0	10^{-6}	—	13.4	0.03
14bc	480	1440	960	17k	10^{-6}	1.0	10^{-9}	—	66.5	0.49
14d	1920	5760	3840	69k	10^{-6}	1.0	10^{-9}	—	67.8	7.07
14e	7680	23040	15360	23k	10^{-1}	1.0	—	—	73.5	6.65
14i	480	1440	960	17k	10^{-6}	1.0	10^{-9}	—	79.7	0.06
14h	1920	5760	3840	69k	10^{-6}	1.0	10^{-9}	—	61.0	0.97
18	1247	3598	2352	25k	10^{-6}	1.0	10^{-6}	0.2	11.1	0.21

(Code in C++, on Intel® Core™ i7-7500U CPU with 2.70GHz 16 GB RAM)

Fig. 20. Statistics for mesh design examples, similar to Fig. 19. w_{prox} is the weight of soft constraints of vertices' proximity to a reference surface.

is expected to be combined with our work on principal-symmetric meshes.

The question of meshing with fair mesh polylines and small variations of normals is not yet exhausted, especially further study of minimizers in the triangle mesh case is needed. Likewise the total absolute curvature deserves a more thorough investigation, also from the mathematical viewpoint.

ACKNOWLEDGMENTS

This work was supported by the SFB-Transregio programme *Discretization in geometry and dynamics*, through grant I2978 of the Austrian Science Fund. The authors are grateful to Felix Günther for useful discussions, and to the anonymous referees for extensive reviews.

REFERENCES

- Lyuba Alboul and Ruud van Damme. 1996. Polyhedral metrics in surface reconstruction. In *The Mathematics of Surfaces VI*, G. Mullineux (Ed.). Clarendon Press, 171–200.
- Ben Andrews. 1999. Gauss curvature flow: the fate of the rolling stones. *Invent. math.* 138 (1999), 151–161.
- Ulrich Bauer, Konrad Polthier, and Max Wardetzky. 2010. Uniform convergence of discrete curvatures from nets of curvature lines. *Discr. Comp. Geom.* 43 (2010), 798–823.
- Ronny Bergmann, Marc Herrmann, Roland Herzog, Stephan Schmidt, and José Vidal Núñez. 2019. Total variation of the normal vector field as shape prior with applications in geometric inverse problems. preprint. <http://arxiv.org/abs/1902.07240>.
- Alexander Bobenko and Yuri Suris. 2009. *Discrete differential geometry: Integrable Structure*. American Math. Soc.
- Alexander I. Bobenko and Peter Schröder. 2005. Discrete Willmore flow. In *Geometry Processing 2005 (Symposium Proceedings)*. 101–110.
- David Bommes, Henrik Zimmer, and Leif Kobbelt. 2009. Mixed-integer quadrangulation. *ACM Trans. Graph.* 28, 3 (2009), 77:1–77:10.
- David Cohen-Steiner and Jean-Marie Morvan. 2003. Restricted Delaunay triangulations and normal cycle. In *Proc. 19th SoCG*. 237–246.
- Manfredo P. do Carmo. 1976. *Differential geometry of curves and surfaces*. Prentice-Hall.
- Nira Dyn, Kai Hormann, Sun-Jeong Kim, and David Levin. 2001. Optimizing 3D triangulations using discrete curvature analysis. In *Mathematical Methods for Curves and Surfaces: Oslo 2000*. Vanderbilt University Press, 135–146.
- Vladimir A. Garanzha. 2010. Discrete extrinsic curvatures and approximation of surfaces by polar polyhedra. *Comp. Math. & Math. Physics* 50 (2010), 65–92.
- Felix Günther, Caigui Jiang, and Helmut Pottmann. 2017. Smooth polyhedral surfaces. preprint. <http://arxiv.org/abs/1703.05318>.
- Klaus Hildebrandt, Konrad Polthier, and Max Wardetzky. 2006. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometriae Dedicata* 123 (2006), 89–112.
- Gerhard Huisken. 1984. Flow by mean curvature of convex surfaces into spheres. *J. Diff. Geometry* 20 (1984), 237–266.
- Caigui Jiang, Felix Günther, Johannes Wallner, and Helmut Pottmann. 2016. Measuring and controlling fairness of triangulations. In *Advances in Architectural Geometry 2016*. vdf Hochschulverlag, Zürich, 24–39.
- Caigui Jiang, Chengcheng Tang, Amir Vaxman, Peter Wonka, and Helmut Pottmann. 2015. Polyhedral patterns. *ACM Trans. Graph.* 34, 6 (2015), 172:1–172:12.
- Leif Kobbelt, Swen Campagna, Jens Vorsatz, and Hans-Peter Seidel. 1998. Interactive multi-resolution modeling on arbitrary meshes. In *Proc. SIGGRAPH*. 105–114.

- David Orden, Günter Rote, Francisco Santos, Brigitte Servatius, Herman Servatius, and Walter Whiteley. 2004. Non-crossing frameworks with non-crossing reciprocals. *Discr. Comp. Geom.* 32, 4 (2004), 567–600.
- Davide Pellis, Martin Kilian, Johannes Wallner, and Helmut Pottmann. 2017. Material-minimizing forms and structures. *ACM Trans. Graphics* 36, 6 (2017), 173:1–173:12.
- Helmut Pottmann, Michael Eigensatz, Amir Vaxman, and Johannes Wallner. 2015. Architectural geometry. *Computers and Graphics* 47 (2015), 145–164.
- Eike Schling, Martin Kilian, Hui Wang, Jonas Schikore, and Helmut Pottmann. 2018. Design and construction of curved support structures with repetitive parameters. In *Proc. Advances in Architectural Geometry*. Klein Publishing, 140–165.
- Oded Stein, Eitan Grinspun, and Keenan Crane. 2018. Developability of triangle meshes. *ACM Trans. Graph.* 37, 4 (2018), 77:1–77:14.
- Rasmus Tamstorf. 2013. *Derivation of discrete bending forces and their gradients*. Technical Report 07/2013. Disney Research.
- Rasmus Tamstorf and Eitan Grinspun. 2013. Discrete bending forces and their Jacobians. *Graph. Models* 75 (2013), 362–370.
- Chengcheng Tang, Xiang Sun, Alexandra Gomes, Johannes Wallner, and Helmut Pottmann. 2014. Form-finding with polyhedral meshes made simple. *ACM Trans. Graphics* 33, 4 (2014), 70:1–70:9.
- Huanxi Zhao and Guoliang Xu. 2006. Triangular surface mesh fairing via Gaussian curvature flow. *J. Comput. Appl. Math.* 195 (2006), 300–311.

APPENDIX

Proof of Prop. 2. The computation of the energy gradient is lengthy but elementary. Similar computations are contained in [Tamstorf 2013]. To find $\nabla \mathcal{E}_v$ we must consider the effect of moving v . In the notation of Figure 9, this movement of v changes quantities ℓ_e , ℓ_e^* and ℓ_e^* , for all edges e incident with v , and all edges \bar{e} opposite to v :

$$\nabla \mathcal{E}_v = \sum_{e \ni v} (\ell_e^* \nabla \ell_e + \ell_e \nabla \ell_e^*) + \sum_{\bar{e} \text{ opposite } v} \ell_{\bar{e}} \nabla \ell_{\bar{e}}^*.$$

The contribution by $\nabla \ell_e = \frac{1}{\ell_e} (v - w)$ accounts for the 1st sum in (1). Next, we observe $e = vw = f \cap f'$, $f = vww^+$, $f' = vww^-$, so

$$\ell_e^* = \|\mathbf{n}_f - \mathbf{n}_{f'}\|, \quad \mathbf{n}_f = \frac{\tilde{\mathbf{n}}_f}{2|f|}, \quad \tilde{\mathbf{n}}_f = (w^+ - w) \times (v - w).$$

For any face $f \ni v$, we find $2\nabla_v |f|$ by rotating the edge vector opposite to v by 90 degrees, i.e., $2\nabla |f| = \mathbf{n}_f \times (w^+ - w)$. We compute the gradient of \mathbf{n}_f , using the notation a_x for the matrix of the linear mapping $x \mapsto a \times x$, so that $\nabla \tilde{\mathbf{n}}_f = (w^+ - w)_x$. Observing $(a_x)^T = -a_x$, we get

$$\begin{aligned} \nabla \mathbf{n}_f^T &= (2|f|)^{-2} (2|f| \nabla \tilde{\mathbf{n}}_f^T - (2\nabla |f|) \tilde{\mathbf{n}}_f^T) \\ &= -(2|f|)^{-1} ((w^+ - w)_x + ((w^+ - w) \times \mathbf{n}_f) \mathbf{n}_f^T) \\ &= -(2|f|)^{-1} (w^+ - w)_x (\text{id} - \mathbf{n}_f \mathbf{n}_f^T) = -(2|f|)^{-1} (w - w^+)_x \pi_f. \end{aligned}$$

From here we compute $\nabla \ell_e^* = \frac{1}{\ell_e^*} (\nabla \mathbf{n}_f^T - \nabla \mathbf{n}_{f'}^T) (\mathbf{n}_f - \mathbf{n}_{f'})$, using the previous expression for $\nabla \mathbf{n}_f$ and the analogous expression for $\nabla \mathbf{n}_{f'}$. Terms $\pi_f(\mathbf{n}_f)$, $\pi_{f'}(\mathbf{n}_{f'})$ vanish, yielding the 2nd sum in Equ. (1). The 3rd sum is analogous and involves $\nabla \ell_{\bar{e}}^*$; however moving v causes a change only to faces on one side of \bar{e} . Accordingly, only faces $f \ni v$ contribute to the 3rd sum. \square

Preparations for ideal energy-minimal case II. The basis of our computations in case of vanishing ℓ_e^* is the following small result:

PROP. 3. Assume that the vertices of mesh $P = (V, E, F)$ move along paths $v(t)$, $0 \leq t < \epsilon$, and that for all edges $e \in E$ we have $\ell_e \neq 0$ and in case $t > 0$ also $\ell_e^* \neq 0$. Then the energy $\mathcal{E}(t)$ at time t obeys

$$\lim_{t \searrow 0} \frac{1}{t} (\mathcal{E}(t) - \mathcal{E}(0)) = \sum_{v \in V} (\nabla_a \mathcal{E}_v)^T \cdot \mathbf{a}_v, \text{ where } \mathbf{a}_v = \frac{dv}{dt} \text{ and}$$

$\nabla_a \mathcal{E}_v$ arises from $\nabla \mathcal{E}_v$ by taking the limit as appropriate ℓ_e^* 's go to zero. The limit of fractions involving ℓ_e^* in the denominator are computed as follows. With the temporary notation $e = v_1 v_2 = f_1 \cap f_2$, we have

$$\lim_{t \searrow 0} \frac{\pi_{f_1}(\mathbf{n}_{f_2})}{\ell_e^*} = \lim_{t \searrow 0} \frac{\mathbf{n}_{f_2}(t) - \mathbf{n}_{f_1}(t)}{\|\mathbf{n}_{f_2}(t) - \mathbf{n}_{f_1}(t)\|} = \pm \frac{J(v_2 - v_1)}{\ell_e}.$$

The symbol J means a 90° rotation within the plane of faces f_j . The sign depends on the velocity vectors \mathbf{a}_v . It points from f_1 to f_2 (resp. from f_2 to f_1) if, as $t \searrow 0$, the edge e is convex (resp. concave) when we define that \mathbf{n}_{f_i} point to the outside.

PROOF. The proposed equation for $d\mathcal{E}/dt$ holds for all $t > 0$, since we are in the generic case of Prop. 2 then. In order to investigate $t = 0$ we observe how $\pi_{f_1}(\mathbf{n}_{f_2})$ changes as the dihedral angle α_e approaches 0. The projected vector is always orthogonal to the edge (like $\mathbf{n}_{f_1} - \mathbf{n}_{f_2}$) and lies in the face f_1 (in the limit, also $\mathbf{n}_{f_1} - \mathbf{n}_{f_2}$ assumes this property). Its length is $\sin \alpha_e$. Since $\ell_e^* = 2 \sin \frac{\alpha_e}{2}$, the limit of $\pi_{f_1}(\mathbf{n}_{f_2})/\ell_e^*$ is a unit vector orthogonal to e , parallel to the plane of faces f_i . Depending on the sign of α_e , it points from f_1 to f_2 or the other way, just as described in the proposition. \square

Proof concerning ideal situation II. To show (2) we assume $e = v_0 v_2$, $f_1 = v_0 v_1 v_2$, $f_2 = v_0 v_2 v_3$. We use the notation $dv_j/dt = \mathbf{a}_j$ for velocity vectors and $\mathbf{e}_{jk} = v_k - v_j$ for edge vectors. Analogous to the proof of Prop. 2 it is not difficult to see that for $t > 0$, we have an equality of the form

$$\frac{d}{dt} \ell_e \ell_e^* = \left(\frac{\ell_e^*}{\ell_e} \mathbf{e}_{23} + \frac{\ell_e \mathbf{e}_{23}}{2|f_2|} \times \frac{\pi_{f_2}(\mathbf{n}_{f_1})}{\ell_e^*} + \frac{\ell_e \mathbf{e}_{12}}{2|f_1|} \times \frac{\pi_{f_1}(\mathbf{n}_{f_2})}{\ell_e^*} \right)^T \mathbf{a}_0 + \dots$$

(further terms involve $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$). According to Prop. 3,

$$\lim_{t \searrow 0} \frac{d}{dt} \ell_e \ell_e^* = \left(\frac{\mathbf{e}_{23}}{2|f_2|} \times (\pm \mathbf{J} \mathbf{e}_{02}) + \frac{\mathbf{e}_{12}}{2|f_1|} \times (\mp \mathbf{J} \mathbf{e}_{02}) \right)^T \mathbf{a}_0 + \dots$$

The vector $\mathbf{J} \mathbf{e}_{02}$ occurs in every summand. A parallelogram shape implies $|f_1| = |f_2|$, $\mathbf{e}_{03} = \mathbf{e}_{12}$, $\mathbf{e}_{01} = \mathbf{e}_{32}$, and we eventually get

$$\begin{aligned} \lim_{t \searrow 0} \frac{d}{dt} \ell_e \ell_e^* &= (\mathbf{e}_{02} \times (\pm \mathbf{J} \mathbf{e}_{02}))^T (\mathbf{a}_1 + \mathbf{a}_3 - \mathbf{a}_0 - \mathbf{a}_2) / 2|f_1| \\ &= \pm \ell_e^2 \mathbf{n}_{f_1}^T (\mathbf{a}_1 + \mathbf{a}_3 - \mathbf{a}_0 - \mathbf{a}_2) / 2|f_1|. \end{aligned}$$

It is not difficult to see that $\mathbf{n}_{f_1}^T (\mathbf{a}_1 + \mathbf{a}_3 - \mathbf{a}_0 - \mathbf{a}_2) < 0$ (resp. > 0) if the edge e is convex (resp. concave) for $t > 0$. At the same time the sign in front of $\mathbf{J} \mathbf{e}_{02}$ has to be chosen as -1 in the convex case and $+1$ in the concave case. In summary, the contribution of the edge e to the total energy has a positive one-sided derivative. \square




Fig. 21. Surface of minimal total absolute curvature. We also show the value of absolute curvature (center; white = zero) and a principal remeshing with edges color coded according to dihedral angle (right).