EUROGRAPHICS 2006 / E. Gröller and L. Szirmay-Kalos
(Guest Editors)

Volume 25 (2006), Number 3

Computing discrete shape operators on general meshes

Eitan Grinspun
Columbia University
eitan@cs.columbia.edu

Yotam Gingold
New York University
gingold@mrl.nyu.edu

Jason Reisman
New York University
jasonr@mrl.nyu.edu

Denis Zorin
New York University
dzorin@mrl.nyu.edu

Abstract

Discrete curvature and shape operators, which capture complete information about directional curvatures at a point, are essential in a variety of applications: simulation of deformable two-dimensional objects, variational modeling and geometric data processing. In many of these applications, objects are represented by meshes. Currently, a spectrum of approaches for formulating curvature operators for meshes exists, ranging from highly accurate but computationally expensive methods used in engineering applications to efficient but less accurate techniques popular in simulation for computer graphics.

We propose a simple and efficient formulation for the shape operator for variational problems on general meshes, using degrees of freedom associated with normals. On the one hand, it is similar in its simplicity to some of the discrete curvature operators commonly used in graphics; on the other hand, it passes a number of important convergence tests and produces consistent results for different types of meshes and mesh refinement.

1. Introduction

Discrete curvature is a key ingredient in a variety of applications: simulation of deformable two-dimensional objects, variational modeling and geometric data processing. In these applications it is often necessary to approximate a solution of a *continuous problem* involving curvature-based energy or forces. Such energy may either capture the physics of the problem (bending energy for thin deformable objects) or our intuition about desirable behavior (variational approaches to surface modeling or curvature flow smoothing of complex geometry).

Curvature and related surface Laplacian discretizations range from highly accurate but complex and computationally expensive high-order finite elements used in engineering to efficient and simple approximations for meshes used in graphics and interactive geometric modeling. Unfortunately, the latter type of approximations, while essential for interactive applications, lacks the predictability and convergence of more mathematically complex and computationally expensive formulations. The lack of convergence leads to mesh-dependent behavior, visual artifacts and fundamental difficulties with adaptive refinement and remeshing of variationally defined shapes.

In this paper we present a simple and efficient discrete curvature operator exhibiting good convergence properties for general meshes. This operator is closely related to widely used discrete geometric operators on one hand, and classical nonconforming finite elements on the other. The discretization we propose uses mesh normals as additional degrees of freedom. One of the essential observations is that the normal field can be defined by scalars, assigned to edges in a natural and geometrically invariant way, similar to the definition of discrete one-forms and discrete vector fields on surfaces. Our local discrete operator uses a constant number of degrees of freedom independent of mesh connectivity and associated forces, and Hessian matrices have simple closed-form expressions.

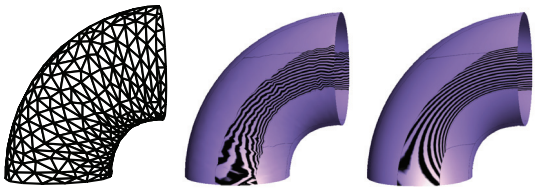

Figure 1: *A mesh (left) optimized using cotangent formula mean curvature discretization (middle) and using midedge normal-based curvature functional (right).*

We present visual and quantitative comparisons of our operator to a representative set of commonly used constructions. Our evaluation is based on practical conditions for convergence, which can be used to evaluate performance of curvature discretizations in variational problems. These conditions are derived from natural geometric invariance properties and mesh independence considerations.

We expect our discretization to be useful in a variety of contexts, including variational modeling and simulation, and in particular for adaptive refinement techniques, which inherently require solutions to be mesh-independent in the limit. The general principles we outline can be applied to designing improved

© The Eurographics Association and Blackwell Publishing 2006. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

higher order discrete geometric functionals, e.g., based on reflection lines or curvature variation.

Terminology. Before proceeding, we briefly review essential concepts used throughout the paper.

The shape operator $\Lambda(\mathbf{p})$ at a point $\mathbf{p}$ of a smooth surface maps tangent vectors $\mathbf{t}$ to normal derivatives in the direction $\mathbf{t}$:

$$\Lambda\mathbf{t} = -\partial_{\mathbf{t}}\mathbf{n} \ .$$

The vector $\partial_t \mathbf{n}$ is tangent to the surface, so the shape operator is a linear map on the tangent vectors. The principal curvatures are the eigenvalues of the shape operator, and the principal curvature directions are its eigenvectors. For a surface parametrized locally at the tangent plane, the shape operator matrix defines the best-fit quadratic approximation to the surface. For a small normal deformation of a flat surface, the shape operator approaches the matrix of second derivatives of the displacements (the Hessian).

Curvature and bending energy functionals. While many functionals are used, we consider the most common examples. The simplest one is the mean curvature-squared (often referred to as Willmore) energy

$$\int_S H^2 dA \ , \tag{1}$$

where H is the mean curvature, and the integral is computed over a surface S. If we consider a deformation g of a surface S, the deformation energy depends on the undeformed shape operator of a surface Λ and the deformed shape operator $\tilde{\Lambda}$. The *bending energy* for thin deformable objects is captured well by the difference of the shape operators for deformed and undeformed surfaces:

$$\int_S \alpha(\mathrm{Tr}\Delta\Lambda)^2 + \beta\mathrm{Tr}(\Delta\Lambda)^2 dA \tag{2}$$

where $\Delta\Lambda = \tilde{\Lambda}\circ g - \Lambda$. This is a variant of the well-known Koiter's model for thin shell deformation [Koi66].

Anisotropic material properties lead to more complex expressions, but in all cases these expressions depend on the surface geometry through the entries of the shape operator.

For many applications (geometric modeling and small deformations), quadratic curvature-based functionals are used. The simplest quadratic energy of this type is

$$\int_D (\Delta\mathbf{f})^2 dA \ , \tag{3}$$

where D is the parametric domain of the surface, $\mathbf{f} : D \to \mathbf{R}^3$, and Δ is the Laplacian. In a linearized setting, the shape operator becomes the matrix of second derivatives, and (3), restricted to normal displacements, corresponds to (1).

General triangle meshes are meshes with no direct restrictions on connectivity (it is common, however, to impose requirements on triangle shapes). *Semi-regular meshes* are meshes obtained by refining a general coarse mesh using a fixed rule which only creates vertices of fixed valence or several valences (e.g., 6, or 4 and 8, depending on the refinement type).

Degrees of freedom is the general term denoting all quantities associated with a mesh which are used to define the energy or PDE discretization; most commonly, these are mesh vertices and surface normal rotations.

Discretization stencil refers to the part of the mesh containing all degrees of freedom used to discretize a single energy or force term, such as two triangles sharing an edge. The total energy or the force vector is obtained by adding these terms together. We are primarily interested in small stencils, especially stencils with a fixed number of degrees of freedom (the simplest example is a triangle). Specific stencils are shown in Figure 2.

Consistency and convergence. While it is hard to expect the point positions obtained by solving a discrete optimization problem to be entirely independent of the choice of mesh, it is reasonable to require that these dependencies vanish for sufficiently fine meshes. Here, we distinguish between *consistency* and *convergence* of curvature estimates, both of which are related to this requirement. We say that a discrete shape operator or curvature is *consistent* if, as we sample a fixed smooth surface increasingly densely, the error between the values of the discrete and continuous shape operators vanishes. We say it is *convergent* if the discrete surface that minimizes a discrete curvature-based energy approaches the solution of the corresponding continuous optimization problem in the limit. To check consistency, we sample a fixed analytically defined surface with increasing density and compare the discrete operator value with the analytic value. To check convergence, we solve a discrete problem on a sequence of increasingly fine meshes with boundary conditions sampled from a known analytic solution to a continuous problem, and we compare the discrete solution to the analytic.

For optimization problems minimizing an energy functional, we use a relatively weak notion of convergence, which we believe to be essential for most applications: as a mesh is refined, we require only *the energy* of the discrete problem's solution to approach the energy of the continuous problem's solution.

2. Previous work

We cannot do justice to the enormous variety of curvature-related discretizations developed in engineering, geometric modeling and computer graphics. In this brief overview of the related work, we focus on approaches most popular in different areas.

The most principled approach for general meshes is to use C^1 finite elements (see e.g., [ZT89] for an overview) or splines, which interpolate the mesh using high-order polynomials, common in engineering, shell simulation, and geometric modeling with splines. Examples of work in variational modeling using these approaches include [CG91, WW92, Gre94]. Subdivision surfaces were used for a similar purpose, e.g., in [MQV97] and [COS00]. For C^1 finite elements, general finite element theory ensures convergence at least for linear problems. These approaches are relatively expensive computationally.

Our focus is closer to the opposite category of methods, which maximize computational efficiency. Each curvature estimate involved in such discretizations depends only on a small number of mesh vertices. We survey various techniques grouped by stencil shape (Figure 2).

Edge and hinge stencils. The simplest stencils are spring (two-point) and hinge (four-point) stencils used for bending in cloth simulation, e.g., [BW98, GDHS03, BMF03]. These approaches, while being surprisingly accurate in many cases, cannot capture the complete local curvature behavior (e.g., principal curvature directions or even the mean curvature value), and in general,

© The Eurographics Association and Blackwell Publishing 2006.

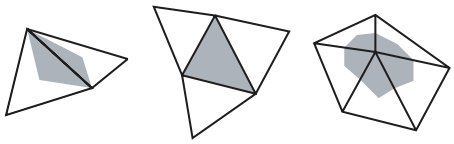

Figure 2: *Commonly used discretization stencils; the area associated with an elementary energy term is indicated in gray.*

resulting forces have a strong mesh dependence. Our technique is more accurate and less mesh-dependent.

Vertex ring stencils. Another common class of methods, used in many geometric modeling applications, is based on local quadratic or cubic surface approximations, or normal curvature fits, typically using a vertex neighborhood [Tau95, GI04]; these discretizations were used in geometric modeling in [MS92, SK01] [WW94] and in engineering, [NU72]. Most of these approaches behave well when used to estimate curvature for sampled surfaces, i.e. they are consistent. As we discuss in more detail below, consistency does not ensure convergence.

Most recently, *discrete geometry* ideas were used to derive a number of valuable curvature operators [PP93, MDSB03]. These methods are robust and have a number of attractive properties but are consistent only for a limited class of meshes [Xu04a, HPW05]. Similar single-ring discretizations, derived from entirely different considerations were also proposed in the engineering literature [OnC93], [OnCRM96]. In [HP04] a vertex-based shape operator based on discrete geometric ideas is introduced.

Triangles with additional degrees of freedom and flaps. Classic finite element constructions use only degrees of freedom assigned to points (nodes) inside the elements (triangles).Notable exceptions are the constructions of [HTC92] and [OnC93], which have the vertex-ring version and the triangle-with-flaps version. In the special case of linear functionals and flat undeformed states, our construction is closely related to the Morley finite element [Mor71]. A more complex DKT element [BBH80] is one of the most commonly used bending elements. In this group, almost all constructions have well-established theoretical convergence properties in the linear case. The discretization we propose belongs in this group. In graphics literature, [Rus04] used a triangle-based construction and vertex normals for curvature estimation.

The idea of using normals as independent degrees of freedom for surface fairing is seen in [Tau01], [YOB02] and [TWBO03]. Most recently, [CDD*04] used mean curvature normals as independent variables in a finite element setting to discretize the Euler-Lagrange equation for the Willmore energy.

[Xu04a] made important advances in understanding the behavior of curvature operators and proposed several new operators. We note that [KCVS98, BK04, CDD*04], rather than solving an energy minimization problem, discretize certain curvature-related PDEs directly. Similar requirements apply to these discretizations; in this paper we restrict our attention to the energy formulations, as the basic principles are easier to understand in this setting.

The convergence conditions we describe are closely related to the well-known *patch test* [IL83] used to evaluate finite elements. We refer to [Stu79, Wan01, ZT97] for the mathematical details of this test. Our preliminary results for this project were presented by the last author at SMI [Zor].

Representative operators. We will use several representative examples of curvature discretizations for comparisons; we briefly describe each of these and refer the reader to the cited papers for details. We do *not* include higher-order operator discretizations, as our goal is to achieve convergence while maintaining the same efficiency of commonly used simple discretizations.

In these definitions and throughout the paper we use $\otimes$ to denote the outer product of two vectors, the linear operator $A = \mathbf{v} \otimes \mathbf{w}$ that maps a vector $\mathbf{x}$ to a vector along $\mathbf{v}$: $A(\mathbf{x}) = (\mathbf{w} \cdot \mathbf{x})\mathbf{v}$. In matrix notation $\mathbf{v} \otimes \mathbf{w} = \mathbf{v}\mathbf{w}^T$.

We use notation for several vectors associated with a triangle depicted in Figure 3.

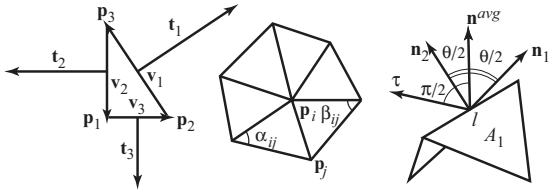

Figure 3: *Left: notation for triangle vectors. Middle: cotangent formula notation. Right: vectors and angles associated with an edge; $\tau = \mathbf{n}^{avg} \times \mathbf{v}$, where $\mathbf{v}$ is the vector along the edge, and the direction of $\mathbf{v}$ is arbitrarily fixed for each edge.*

Hinge energy and operator. This energy is widely used for cloth simulation, e.g., [BW98, GDHS03, BMF03]. It can be considered an approximation of the Willmore energy. Its stencil (Figure 2, left) consists of two triangles, and in one possible formulation, the energy for an edge is computed as

$$\frac{3(\theta-\tilde{\theta})l^2}{A_1+A_2},$$

where θ and $\tilde{\theta}$ are the angles between two normals, in deformed and undeformed states respectively, l is the edge length, and A_i are the areas of two triangles. The associated shape operator is naturally defined as $\theta\tau \otimes \tau/(A_1+A_2)$.

Quadratic fit vertex operator. This is a common approach to curvature estimation; we use the normal curvature fit approach similar to [MS92, Tau95, SK01]. As the construction is relatively lengthy, we do not provide explicit formulas. The idea is to construct a circle for each edge at a vertex and regard it as the normal curvature approximation, fitting a quadratic function that minimizes least-squares error in normal curvatures.

Cotangent formula vertex energy. This widely used energy, which has a vertex one-ring stencil, approximates the Willmore energy. Per vertex it is defined as

$$\left[\frac{1}{4A_i}\sum_{j\in N(i)}(\cot\alpha_{ij}+\cot\beta_{ij})(\mathbf{p}_j-\mathbf{p}_i)\right]^2,$$

where $N(i)$ is the set of indices of vertices adjacent to i and the vertex area A_i can be computed in several ways as discussed in [MDSB03]; we found the A_{mixed} formulation most accurate, but results are similar for one third of the area of triangles in $N(i)$. (see Figure 3 for other notation). A compatible shape operator is described in [HP04].

© The Eurographics Association and Blackwell Publishing 2006.

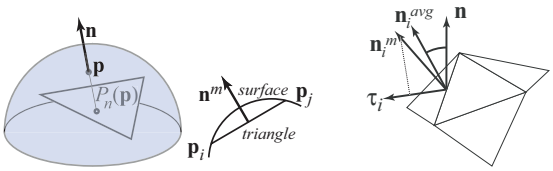

Figure 4: *Left: Normal map from the surface to the mesh. Middle: The midedge normal is up to $O(h^2)$ perpendicular to the edge.*

Triangle-averaged operator. This operator, to the best of our knowledge, was not described in exactly this form, although it is highly similar to the finite element discretization of [OnC93]. Its stencil is a triangle with flaps, and the expression for the shape operator is

$$\sum_{i=1,2,3} \frac{\theta_i}{2Al_i} \mathbf{t}_i \otimes \mathbf{t}_i \ , \tag{4}$$

where the summation is over three triangle edges, θ_i is the angle between triangle normals for two triangles meeting at the edge i, and $\mathbf{t}_i$ are vectors of length l_i perpendicular to triangle edges. This operator is similar to [HP04], but the averaging is over three edges of a triangle, rather than all edges adjacent to a vertex, and no projection to the tangent plane of a triangle is necessary.

This operator is probably the simplest operator using only positional degrees of freedom and capable of reproducing arbitrary curvature directions. The operator we propose can be regarded as a corrected version of this operator, which is why we consider it for comparison purposes.

3. Midedge normal shape operator

The need for better, more efficient constructions can be seen in Figure 12. For a simple deformation, some discretizations produce results that depend heavily on the chosen mesh while others converge to a mesh-independent value.

In this section we derive our construction. In the derivation, we only aim to ensure consistency and some natural properties for the discrete shape operator. However, the construction also happens to satisfy important convergence conditions.

For simplicity of exposition, we consider a closed surface, but the construction applies in the case of surfaces with boundaries with minimal changes.

Let S be our surface and S_h a mesh approximating the surface with vertices on the surface.

For a sufficiently fine mesh, the normal map $P_n : S \to S_h$ maps every point of the surface $\mathbf{p}$ to the point of the mesh S_h obtained by intersecting the mesh with the line passing through $\mathbf{p}$ in direction $\mathbf{n}$ (Figure 4). This map is one-to-one for sufficiently fine meshes S_h.

We make the following simple observation. Consider the midpoint $\mathbf{p}^m$ of an edge $e = (\mathbf{p}_1, \mathbf{p}_2)$, and the corresponding point on the surface $\tilde{\mathbf{p}}^m = P_n^{-1}(\mathbf{p}^m)$. Figure 4 shows the cross-section of the surface by the plane spanned by the edge e and the normal $\mathbf{n}(\tilde{\mathbf{p}}^m)$. If we approximate this cross-section by a quadratic curve interpolating $\mathbf{p}_1$ and $\mathbf{p}_2$, then the normal $\mathbf{n}^m = \mathbf{n}(P_n^{-1}(p^m))$ (the midedge normal) is perpendicular to the edge e. Thus, for curvature approximation of the lowest order, one can fully represent the midedge normal by its angle of rotation around the edge or another equivalent scalar variable. This representation plays an important role in our construction, and can also be considered as a natural definition of a discrete normal. This concept resembles the idea of representing discrete 1-forms and vector fields by scalars associated with edges, e.g., [GY02]. The connection can be made more precise if we observe that the normals near a point $\mathbf{p}$ can be represented by their projections on the tangent plane at $\mathbf{p}$. In particular, we can use the plane of the triangle parallel to the tangent plane at a point on the surface.

For normals perpendicular to edges, the projected vectors will also be perpendicular to edges. By analogy with the continuous shape operator, one can now define a per-triangle discrete shape operator as the gradient of the normal over the triangle, assuming the normal projection varies linearly on the triangle.

Based on this simple definition, we can easily compute an invariant representation of the shape operator in terms of triangle edges and midedge normals as explained below. The notation we use for vectors associated with a triangle is shown in Figure 3 (left). We use $\widehat{\mathbf{v}}$ to denote a unit-length vector along $\mathbf{v}$. We use i, j, k to denote a cyclic permutation of indices $1, 2, 3$.

First, we specify the midedge normal in a frame symmetric (up to a sign) with respect to both triangles sharing the normal. Let $\mathbf{n}_i^{avg}$ be the average unit normal of the triangle T, and $\tau_i^{avg} = \mathbf{v}_i \times \mathbf{n}_i^{avg}$. The vector τ_i obtained from the other triangle sharing the edge i has the opposite sign; we arbitrarily choose one of the triangles as the owner of the edge and use its τ_i. In this basis, we define the midedge normal as the unit vector along $\varphi_i \tau_i + \mathbf{n}_i^{avg}$ where φ_i is the new variable we introduce. (A somewhat more natural definition of φ_i would be the angle between the midedge normal and $\mathbf{n}_i^{avg}$; however, these definitions do not differ in the limit of small angles, and this definition simplifies the derivation.)

We define our shape operator by its action on tangent directions, using a finite difference version of the continuous definition:

$$\Lambda \frac{1}{2} \mathbf{v}_i = \mathbf{n}_j^m - \mathbf{n}_k^m \ , \tag{5}$$

where we assume the operator is constant on the triangle, and the tangent plane to the surface is approximated by the plane of the triangle. These equations need to be solved for components of Λ.

One can observe that this system may have no solution if the differences of midedge normals are not parallel to the plane. To obtain a valid system, we perturb the normals by an amount quadratic in triangle size, which does not change the approximation order of the operator. Specifically, we substitute for the normals in equation 5 the expressions $\mathbf{n} + (s_l \varphi_l + \theta_l/2)\widehat{\mathbf{t}_l}$, $l = j, k$, where θ_l is the angle between normals of the triangles T and T_l, and s_l is 1 or -1 depending on the ownership of the normal. As a result, $\mathbf{n}$ vanishes, and the resulting vectors are in the triangle plane.

In this case, a simple invariant expression for the discrete shape operator can be obtained by using $\mathbf{t}_i \otimes \mathbf{t}_i$, $i = 1, 2, 3$ as a basis. As each of these operators is symmetric, the resulting operator is guaranteed to be symmetric.

© The Eurographics Association and Blackwell Publishing 2006.

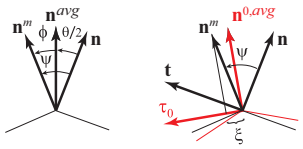

Figure 5: *Variables used in the midedge operator construction and the change of variables.*

Substituting this general form into 5 and solving for the coefficients of three basis operators, we obtain (Figure 5)

$$\Lambda = \sum_{i=1,2,3} \frac{\theta_i/2 + s_i\varphi_i}{Al_i}\mathbf{t}_i \otimes \mathbf{t}_i = \sum_{i=1,2,3} \frac{\psi_i}{Al_i}\mathbf{t}_i \otimes \mathbf{t}_i \,, \qquad (6)$$

where A is the are of the triangle, and $\psi_i = \theta_i/2 + s_i\varphi_i$. We note that this operator depends on three triangle vertices, midedge scalar variables φ_i and vertices of adjacent triangles because the θ_i depend on them.

We make the following observation: this operator is a modification of the triangle average operator, with additional edge degrees of freedom used to "correct" the angles between average normals and triangle normals.

Next, for computational purposes, we perform a change of variables, reducing to a form which does not depend on adjacent triangle vertex positions.

Change of variables. While the formulation above has a clear geometric meaning, for efficient computation it is desirable to reduce the number of degrees of freedom on which each term depends. Note that choosing the average normal as a reference for defining the midedge normal orientation is not essential; it is simply most natural. Indeed, by (6) the shape operator may be expressed in terms of the total angle, ψ_i, between the triangle normal and the midedge normal.

For efficient computation, one may use *fixed* spatial directions a_i in the world coordinate system as references. If this is done, there is no longer any dependence on vertices of adjacent triangles; but one needs to take into account the rotation of the triangle with respect to the global coordinate system.

The problem with using a fixed global system (known as the "total Lagrangian approach") is that the angles quickly become large, and an edge of the triangle can become parallel to the corresponding global axis a; in this case, the midedge normal can no longer be reconstructed from its projection onto a.

An alternative solution is to have the reference axes follow the object ("updated Lagrangian approach"). Fix a time t_0 and denote the corresponding vectors $\boldsymbol{\tau}_i(t_0)$ as $\boldsymbol{\tau}_i^0$. We introduce a new variable ξ_i, defined as the projection of the midedge normal on $\boldsymbol{\tau}_i^0$, choosing sign according to ownership. Note that the definition of ξ_i is symmetric with respect to permutation of the two triangles, up to the sign determined by ownership.

As $\mathbf{n}_i^m \approx \mathbf{n} + \psi_i\mathbf{t}_i$, and by the definition $\xi_i = (\mathbf{n}_i^m \cdot \boldsymbol{\tau}_i^0)$, we can compute ψ_i for the shape operator expression (6) as a function of ξ_i. This leads to the following discrete shape operator depending on ξ_i. This new operator is not identical to the one we started with but approximates it to the second order:

$$\Lambda = \sum_{i=1,2,3} \frac{s_i\xi_i - (\mathbf{n}\cdot\boldsymbol{\tau}_i^0)}{Al_i(\widehat{\mathbf{t}_i}\cdot\boldsymbol{\tau}_i^0)}\mathbf{t}_i \otimes \mathbf{t}_i \,, \qquad (7)$$

Both forms of our shape operator can be easily shown to be consistent, under the standard assumption that the minimal angle of any refined mesh triangle is bounded from below by a constant. What sets it aside from such consistent operators as local quadratic fit, is that in addition to consistency, this operator has remarkable convergence properties, as it satisfies convergence conditions other operators failed to meet; these properties are confirmed by a number of experiments in Section 5.

Numerical solvers typically require implementation of first and second derivatives of energy. In the Appendix we describe a simplified treatment of these derivatives for energies discretized using (7).

4. Evaluation of discrete shape and curvature operators

If a discretization is derived based on informal considerations, which is often the case for the simplest and most efficient approaches, it may be hard to predict its behavior and its degree of mesh dependence. It is useful to have a set of easy to check desirable conditions.

The geometric invariance conditions are the easiest to check and are typically satisfied by most geometrically constructed operators, including all in our representative set. The essential conditions include the following: *the energy is invariant with respect to rigid transformations, the energy is invariant with respect to uniform scale* (a general property of functionals which are integrals of expressions quadratic in curvature), and *if both deformed and undeformed configurations are planar, the energy vanishes* (planar surfaces have zero curvatures).

It is desirable for a shape operator and the associated energy to be consistent as described in the introduction. It is easy to see why this requirement is important from the physical point of view: if it is not satisfied, we cannot reliably estimate the shape operator of the surface from its discrete approximation. In the case of physically-based simulation, depending on the choice of the sequence of refined meshes, different deformations can be obtained from the same external forces acting on the surface. Yet several commonly used discretizations fail consistency tests for many meshes.

Convergence. Our discussion here is informal. A more formal derivation of these conditions, extending the patch test conditions, will be presented in a separate report.

When discussing convergence, we need to consider sequences of meshes with triangle size going to zero. Usually, additional conditions are imposed to ensure convergence: a weak constraint may be that the triangle aspect ratio remains bounded, while a stronger constraint may require a sequence of regularly subdivided meshes. We call mesh sequences satisfying these constraints *admissible*. Unfortunately, consistency does not guarantee convergence. This can be seen from the following simple example, illustrated in Figure 6 in one dimension. If one takes a consistent discretization of the curve energy $\int (f'')^2 dx$, $\sum_i (f_{i+1} - 2f_i + f_{i-1})/h^2$, and removes all odd terms, multiplying the remaining terms by two, the discretization remains consistent; however, for any boundary conditions, the minimal energy configuration will be zero,

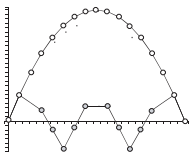

Figure 6: *Minimizing consistent, but not convergent discrete energy (lower curve; the upper curve is the correct solution).*

© The Eurographics Association and Blackwell Publishing 2006.

as one can simply set $f_{2i+1} = (f_{2i} + f_{2i+2})/2$. This is clearly not true for the solution of the continuous problem with fixed tangents and positions on the boundary.

In general, formulating exact convergence conditions for displacements would require introducing bases for the surface discretizations and suitable functional norms.

Fortunately, if one only wants necessary conditions, one can consider convergence of the discrete energy of the approximate solutions to the continuous energy of the exact solution. While this appears to be a relatively weak requirement, not necessarily implying convergence (e.g., pointwise or average), it turns out that it leads to quite strong and explicit tests easily applied to evaluating discretizations. We formulate the general requirement more precisely before explaining the practical tests based on this requirement.

Energy convergence. *Suppose for an energy minimization problem, certain boundary conditions guarantee the solution is unique. Then, for any sequence of admissible meshes with triangle size going to zero, and discrete boundary conditions sampled from the continuous solution, the energy of the solution of the corresponding discrete problem converges to the energy of the solution of the continuous problem.*

Convergence conditions for linearized problems. To convert this general condition to tests, we consider a very specific case: small normal deformations of planar surfaces (plates) which can be expressed as quadratic functions on the plane of the undeformed configuration. Any "good" general discrete energy function should also work in this simple case. Any functional of the form (2) in linearized form reduces to

$$\int_D \alpha(z_{uu} + z_{vv})^2 + \beta(z_{uu}^2 + 2z_{uv}^2 + z_{vv}^2) dA \ , \qquad (8)$$

where z is the scalar displacement perpendicular to the plate. It is easy to show that any quadratic function is a stationary point of this functional. For suitable boundary conditions (positions and normal derivatives on the boundary), minimization of the functional recovers quadratic functions exactly from their boundary values.

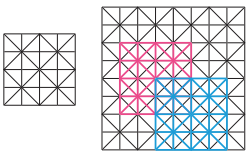

Figure 7: *A tilable patch.*

We consider a general type of *local discretization*: we assume the mesh is partitioned into subdomains, which overlap only at boundaries, and the curvature energy for the whole surface is computed as the sum of total energies for each subdomain. Example subdomains are shown in Figure 2. The energy on each subdomain is computed locally, that is, only using vertices which are contained within certain mesh neighborhoods of the domain: $E(p) = \sum_i E(p_i)$, where the summation is over all subdomains, p is the vector of degrees of freedom, and p_i are the vectors of degrees of freedom for the i-th subdomain. We will show that a necessary condition for energy convergence is that the discrete functional recovers the samples of a quadratic function q *exactly* on a certain class of small meshes (*tilable patches*) if the boundary conditions are sampled from q. We call a planar mesh P a tilable patch if it is possible to construct an admissible sequence of refined meshes $\{M_n\}$, with each M_n containing at least n patches identical to P, such that the fraction of area of M_n covered by these patches remains bounded below by a constant c as $n \to \infty$. In addition, we require P to contain at least one free degree of freedom once discrete Dirichlet and Neumann boundary conditions are specified.

Explicit condition on linearized discretizations.

Proposition 1

Consistency test. *If the displacements z for a planar mesh are sampled from a quadratic function on a tilable patch, the discrete energy density on the patch should match the continuous energy of the quadratic function exactly.*

Energy convergence test. *If the discrete energy density is constant for a polygonal mesh, and its sampling pattern can be tiled, the total energy should be minimal with respect to variation of all vertices of the patch not fixed by discretization of position and normal derivative (Dirichlet and Neumann) boundary conditions.*

There are two remarkable features of this test: (1) one needs to check the conditions for small fixed-size meshes which can be tiled in the sense described above. For example, for all stencils shown in Figure 2, it is sufficient to consider a double ring of vertices around a vertex; (2) only an exact reproduction of three quadratic functions is necessary, assuming that the energy for the linear functions is zero.

We note that the variety of patches which needs to be checked strongly depends on the way admissible refinements are defined: for example, if only uniformly refined meshes of the type used in subdivision are allowable, one only needs to check patches with regular connectivity.

For small normal displacement, u_z, for a flat plate,

$$E(p^0 + u_z) \approx \frac{1}{2} z^T H_z(p^0) z \ ,$$

where $H_z(p^0)$ is the matrix of second derivatives of the energy at p^0 with respect to scalar normal displacements and z is the vector of scalar normal displacements. Using scale invariance of $E(p)$, we easily get the following scaling property: $H_z(sp_0) = (1/s^2) H_z(p_0)$.

Building on this scaling property, we provide some intuition for the convergence conditions. Suppose on a tilable patch P a discretization has a nonzero energy error for a quadratic function. We can shrink the patch by a factor h^2 and create a mesh for which the number of patches identical to P increases as h^2, so the fraction of the area these patches occupy remains bounded from below. As we have observed with rescaling of the domain, the linearized energy for each patch scales as h^2; as a consequence, the energy error also scales as h^2. On the other hand, as the number of patches identical to $h^2 P$, with error $h^2\varepsilon$ increases as h^2, the error remains constant as the mesh is refined. The second condition is justified in a similar way.

This proposition closely parallels the well-known patch test widely used in the finite element community, but does not assume an underlying basis.

Properties of representative operators. We summarize the properties of the linearized versions of the operators described in Section 2. The *hinge operator* is only consistent and convergent on equilateral meshes (after rescaling by a factor of 6) and when used to compute the mean curvature. The *vertex normal fit* operator is consistent for vertices of valence over five (for low-valence vertices, one in principle needs to consider more than a

© The Eurographics Association and Blackwell Publishing 2006.

single-ring neighborhood). It does not pass the quadratic reproduction test for many mesh refinement sequences as can be seen from the results in Section 5. The *cotangent formula* discretization of the mean curvature energy is consistent and convergent on regularly refined meshes, affine transformations of regularly refined meshes, and meshes obtained from regular meshes by edge flips (but not combined with affine transformations). The *triangle averaged* operator is consistent and convergent for regular meshes and equilateral meshes but not for general mesh refinement. The *midedge normal* operator is consistent and convergent for general meshes, for which the triangle aspect ratio remains bounded under refinement. We observe that due to its small stencil, any ring of triangles is a tilable patch for this operator. The exact quadratic reproduction can be verified both numerically (see Figure 14, left) and analytically. This property can be either obtained indirectly, by proving equivalence to the Morley element in the plate case, or by direct geometric check for single-ring patches.

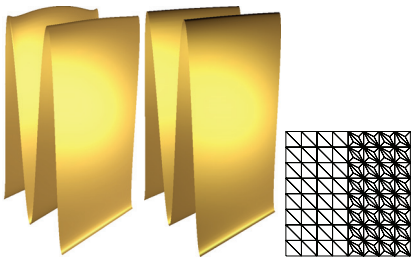

Figure 8: *The analytic solution for this problem setup is a function of y only; Left: results for the cotangent formula; middle: results for the midedge normal operator. For the regular mesh, both operators yield close approximations of the correct solution. Mesh structure is shown on the right.*

5. Numerical comparisons

We focus on testing our shape operator on examples for which the answers are known and can be computed quantitatively, or the quality of results can be easily estimated visually.

Figure 10, compares estimated curvature to analytically computed values, demonstrating the consistency of our operator. Note that our operator uses normals to estimate the curvature; unlike the other two operators, it cannot be used to directly estimate mesh curvature, unless the normals are known (e.g., acquired as a part of the 3d scanning process or computed from the mesh by other means).

For a number of problem types, we compare performance of our operator with other available formulations. We acknowledge that the convergence rate is better for many classical finite elements and high order surfaces, however they carry the burden of increased computational cost and implementation complexity. In contrast, our goal is to achieve convergence and mesh-independent behavior without sacrificing simplicity and efficiency; therefore, we focus on comparisons to discretizations widely used in computer graphics and geometric modeling with meshes.

First, we compare the behavior of our operator and several other formulations for a simple bending problem, for which exact solutions are known: a square plate with a uniform unit load and fixed boundaries (Figure 12).

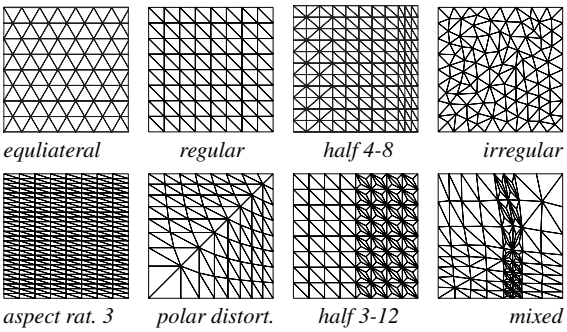

Figure 9: *Mesh types used in our tests*

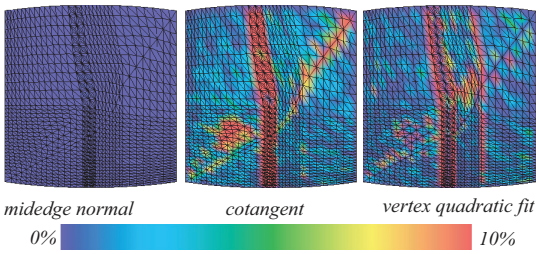

Figure 10: *Mean curvature error on a unit cylinder for three operators. Errors above 10% are truncated.*

We use several types of meshes, shown in Figure 9. We observe that the cotangent operator and our operator are the only ones with consistently convergent results.

We verify that our operator has the expected accurate behavior in the nonlinear case by minimizing the Willmore energy (equivalent to the integral of H^2 under appropriate boundary conditions) and comparing the result to explicitly known solutions for spheres and circular cylinders. In the latter case, an additional area term is needed in the energy to obtain a cylinder of a fixed radius. As the plots in Figure 12 show, minimizing our energy discretization recovers these shapes exactly; the behavior for the cylinder is similar.

Good displacement convergence, while essential in many cases, is insufficient for high-quality surface generation. For example, if we need to model deformations of highly reflective surfaces, the quality of the result will be dependent on the behavior of the surface normals. The next two sets of examples show shapes and their reflection lines obtained for fixed cylindric boundary conditions using different operators. The examples in Figure 14 show the behavior of different operators; the quadratic examples on the left simultaneously show the degree of deviation from the convergence condition. Observe that our operator produces the exact result for all meshes.

Next, we consider anisotropic nonflat undeformed shapes, for which the deformation energy cannot be captured by the mean curvature energy alone and requires the full curvature operator (see Figure 13). In addition, the energy includes an in-plane stretching component. We use the well-established engineering finite-element obstacle course examples [MH85], for which the correct solutions are known for linearized functionals. Our operator yields convergent results for these examples. Once more a typical shell finite element on a coarse mesh would

© The Eurographics Association and Blackwell Publishing 2006.

yield higher accuracy, and incur a higher computational expense.

Using normals as degrees of freedom (DOFs) has the advantage of yielding optimized surface normals directly. Figure 11 compares a surface with half-3-12 mesh connectivity optimized using the cotangent formula, the midedge normal operator with normals computed by averaging face normals, and the midedge normal operator with normals computed from edge DOFs.

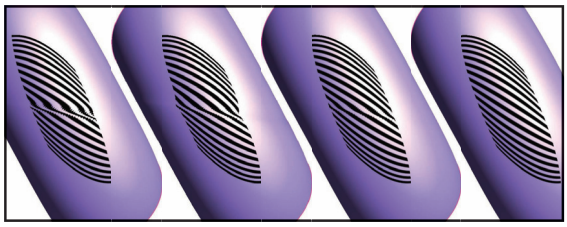

Figure 11: *From left to right: a surface obtained by minimizing thin plate energy using the cotangent discretization; same surface after 10 iterations of normal smoothing; a surface obtained using the midedge normal discretization, with vertex normals computed by averaging face normals; same surface with midedge normals used to compute vertex normals.*

Finally, we compare the computational cost of several operators (see Table 1). We consider two measures of efficiency: the number of floating-point operations needed to evaluate the energy or its Hessian, and the number of nonzeros in the Hessian matrix, which typically determines the cost of a solve. Our operator has a distinctive feature: unlike operators that use only vertex positions, its number of DOFs does not coincide with the number of vertices. Depending on the application, it may be relevant to consider work per DOF or per vertex. For problems involving Hessians of the energy, the most relevant factor is the *number of nonzeros* (nnz) in the matrix and the expense of computing the matrix. This cost is approximately proportional to the number of DOFs per elementary energy term, squared.

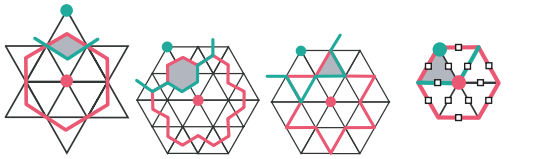

Hinge *Vertex ring* *Triangle+flaps* *Midedge normal*

stencil	linear			nonlinear		
	$\frac{\text{nnz}}{\text{DOF}}$	$\frac{\text{nnz}}{\text{vtx}}$	$\frac{\text{nnz}}{\text{term}}$	$\frac{\text{nnz}}{\text{DOF}}$	$\frac{\text{nnz}}{\text{vtx}}$	$\frac{\text{nnz}}{\text{term}}$
Hinge	13	13	4	39	117	12
Vertex ring	19	19	7	57	171	21
Triangle +flaps	19	19	6	57	171	18
Midedge norm.	11.5	46	6	25	150	12

Table 1: *Comparison of computational cost for a regular grid, for linear and nonlinear problems; typically, the average numbers for an arbitrary grid are close. For the hinge, triangle-averaged and midedge normal stencils the number of DOFs per energy term in the interior of the mesh does not depend on mesh connectivity. The diagrams show in red the boundary of the area in which all DOFs share an energy term with the vertex DOF marked with the red dot. For the midedge normal operator we average over all (vertex and edge) DOFs.*

6. Conclusions and future work

We have described criteria for evaluating shape operators for variational problems and presented a simple shape operator formulation, using normals as degrees of freedom, which bridges the gap between techniques commonly used in computer graphics and geometric modeling, on the one hand, and engineering, on the other.

We have demonstrated that our shape operator has many desirable accuracy and convergence properties yet is easy to implement. We believe the proposed formulation will be useful for a variety of applications in graphics and geometric modeling. We will make the software implementing our discretization as well as other discretizations, to which we are comparing, publicly available.

The principles used in the convergence condition we have described are useful for other types of functionals; for example, curvature variation is an important criterion for shape optimization. Similar tests can be derived for these types of functionals.

While being most natural, normals are not the only possible degrees of freedom one can add to the mesh; e.g., scalar mean curvature may be a useful degree of freedom.

We have observed that midedge normals are represented in a natural way as scalars of edges, similar to 1-forms; our discretization can be viewed as a corrected version. An interesting direction for future research is to explore this connection and consider discrete geometry definitions involving higher order quantities.

In the case of quasi-isometric surface deformations, our shape operator can serve as a starting point for formulating a quadratic bending energy, following the technique described in [BWH*06]. The resulting formulation, having bending forces linear in mesh positions, is suitable for fast simulation of inextensible thin plates.

Appendix: computing derivatives. While it is relatively straightforward to compute complete first and second derivatives of energy discretizations based on our shape operator, we observe that one can successfully use a simplified form, which we applied in all our calculations. Consider for example, the energy $\mathrm{Tr}\Lambda^2$. If the deformation of triangles with respect to the reference configuration is small, in the expression e.g., for $\mathrm{Tr}\Lambda^2 = \sum_{i,j} c_i c_j (s_i\xi_i - f_i)(s_j\xi_j - f_j) T_{ij}$, where $T_{ij} = (\mathbf{t}_i \cdot \mathbf{t}_j)^2$, $c_i = 1/(A l_i (\widehat{\mathbf{t}_i \cdot \boldsymbol{\tau}_i^0}))$, and $f_i = (\mathbf{n} \cdot \boldsymbol{\tau}_i^0)$, the changes in all variables are of higher order compared to f_i. Under these assumptions, only the expressions for derivatives of f_i are needed. The first derivative is particularly simple: $\frac{\partial f_i}{\partial \mathbf{p}_k} = \frac{-(\boldsymbol{\tau}_i \cdot \mathbf{t}_k)\mathbf{n}}{2A}$. The second derivatives are given by

$$\frac{\partial^2 f_i}{\partial \mathbf{p}_{k_1} \partial \mathbf{p}_{k_2}} = -\left(\frac{1}{4A^2} \left((\boldsymbol{\tau}_i^0 \cdot \mathbf{t}_{k_1})(\mathbf{n} \otimes \mathbf{t}_{k_2} + \mathbf{t}_{k_2} \otimes \mathbf{n}) + [\mathbf{t}_{k_2}, \boldsymbol{\tau}_i^0, \mathbf{v}_{k_1}](\mathbf{n} \otimes \mathbf{n}) \right) + \frac{1}{2A} r(k_1, k_2) \mathbf{n} \otimes (\mathbf{n} \times \boldsymbol{\tau}_i^0) \right),$$

where $r(k_1, k_2)$ is defined to be 0 if $k_1 = k_2$, -1 if k_2 immediately follows k_1, and 1 if k_2 immediately precedes k_1.

Acknowledgements. The research presented in this paper was supported by NSF awards CCR-0093390, DMS-0138445, CCF-0528402, and the IBM Faculty Partnership Award. We thank Adrian Secord and Jeff Han for their contributions to the software used in this work.

© The Eurographics Association and Blackwell Publishing 2006.

References

[BBH80] BATOZ J. L., BATHE K. J., HO L. W.: A study of three-node triangular plate bending elements. *International Journal for Numerical Methods in Engineering 15* (1980), 1771.

[BK04] BOTSCH M., KOBBELT L.: An intuitive framework for real-time freeform modeling. *ACM Transactions on Graphics 23*, 3 (2004), 630–634.

[BL89] BATOZ J. L., LARDEUR P.: A discrete shear triangular nine d.o.f. element for the analysis of thick to very thin plates. *International Journal for Numerical Methods in Engineering 29* (1989), 1595.

[BMF03] BRIDSON R., MARINO S., FEDKIW R.: Simulation of clothing with folds and wrinkles. *ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2003), 28–36.

[BW98] BARAFF D., WITKIN A.: Large steps in cloth simulation. *Computer Graphics. Proceedings. SIGGRAPH 98 Conference Proceedings* (1998), 43–54.

[BWH*06] BERGOU M., WARDETZKY M., HARMON D., ZORIN D., GRINSPUN E.: A quadratic bending model for inextensible surfaces. *Europgaphics Symposium on Geometry Processing* (2006).

[CDD*04] CLARENZ U., DIEWALD U., DZIUK G., RUMPF M., RUSU R.: A finite element method for surface restoration with smooth boundary conditions. *Comput. Aided Geom. Design 21*, 5 (2004), 427–445.

[CG91] CELNIKER G., GOSSARD D.: Deformable curve and surface finite-elements for free-form shape design. *Computer Graphics 25*, 4 (1991), 257–266.

[COS00] CIRAK F., ORTIZ M., SCHRÖDER P.: Subdivision surfaces: A new paradigm for thin-shell finite-element analysis. *Internat. J. Numer. Methods Engrg. 47*, 12 (2000), 2039–2072.

[GDHS03] GRINSPUN E., DESBRUN M., HIRANI A., SCHRÖDER P.: Discrete shells. In *Proceedings of ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (2003), Breen D., Lin M., (Eds.).

[GI04] GOLDFEATHER J., INTERRANTE V.: A novel cubic-order algorithm for approximating principal direction vectors. *ACM Transactions on Graphics 23*, 1 (Jan. 2004), 45–63.

[Gre94] GREINER G.: Variational design and fairing of spline surfaces. *Computer Graphics Forum 13*, 3 (1994), 143–154.

[GY02] GU X., YAU S.-T.: Computing conformal structures of surfaces. *Communications in Information and Systems 2*, 2 (2002).

[HP04] HILDEBRANDT K., POLTHIER K.: Anisotropic filtering of non-linear surface features. *Computer Graphics Forum 23*, 3 (2004).

[HPW05] HILDEBRANDT K., POLTHIER K., WARDETZKY M.: On the convergence of metric and geometric properties of polyhedral surfaces. Submitted for publication, 2005.

[HTC92] HAMPSHIRE J. K., TOPPING B. H. V., CHAN H. C.: Three node triangular elements with one degree of freedom per node. *Engineering Computations 9* (1992), 49.

[IL83] IRONS B., LOIKKANEN M.: An engineers' defence of the patch test. *Internat. J. Numer. Methods Engrg. 19*, 9 (1983), 1391–1401.

[KCVS98] KOBBELT L., CAMPAGNA S., VORSATZ J., SEIDEL H. P.: Interactive multi-resolution modeling on arbitrary meshes. *Computer Graphics. Proceedings. SIGGRAPH 98 Conference Proceedings* (1998), 105–114.

[Koi66] KOITER W. T.: On the nonlinear theory of thin elastic shells. I, II, III. *Nederl. Akad. Wetensch. Proc. Ser. B 69* (1966), 1–17, 18–32, 33–54.

[MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and Mathematics III*, Hege H.-C., Polthier K., (Eds.). Springer-Verlag, Heidelberg, 2003, pp. 35–57.

[MH85] MACNEAL R. H., HARDER R. L.: A proposed standard set of problems to test finite element accuracy. *Computer and Structures 20* (1985), 121.

[Mor71] MORLEY L. S. D.: On the constant moment plate bending element. *Journal of Strain Analysis 6* (1971), 10.

[MQV97] MANDAL C., QIN H., VEMURI B. C.:. In *In Proceedings of IEEE Visualization '97* (1997), pp. 371–377.

[MS92] MORETON H. P., SEQUIN C. H.: Functional optimization for fair surface design. *Computer Graphics 26*, 2 (1992), 167–176.

[NU72] NAY R. A., UTKU S.: An alternative to the finite element method. *Variational Methods Engineering 1* (1972).

[OnC93] OÑATE E., CERVERA M.: Derivation of thin plate bending elements with one degree of freedom per node. *Engineering Computations 10* (1993), 543.

[OnCRM96] OÑATE E., CENDOYA P., ROJEK J., MIQUEL J.: Non linear explicit dynamic analysis of shell structures using a simple triangle with translational degrees of freedom only, 1996.

[PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experiment. Math. 2*, 1 (1993), 15–36.

[Rus04] RUSINKIEWICZ S.: Estimating curvatures and their derivatives on triangle meshes. In *Symposium on 3D Data Processing, Visualization, and Transmission* (2004).

[SK01] SCHNEIDER R., KOBBELT L.: Geometric fairing of irregular meshes for free-form surface design. *Computer-Aided Geometric Design 18*, 4 (2001), 359–379.

[Stu79] STUMMEL F.: The generalized patch test. *SIAM J. Numer. Anal. 16*, 3 (1979), 449–471.

[Tau95] TAUBIN G.: Estimating the tensor of curvature of a surface from a polyhedral approximation. *Proceedings. Fifth International Conference on Computer Vision (Cat. No.95CB35744)* (1995), 902–907.

[Tau01] TAUBIN G.: *Linear Anisotropic Mesh Filtering*. Tech. Rep. RC-22213 10/18/2001, IBM T.J.Watson Research Center, 2001.

[TWBO03] TASDIZEN T., WHITAKER R., BURCHARD P., OSHER S.: Geometric surface processing via normal maps. *ACM Trans. Graph. 22*, 4 (2003), 1012–1033.

[Wan01] WANG M.: On the necessity and sufficiency of the patch test for convergence of nonconforming finite elements. *SIAM J. Numer. Anal. 39*, 2 (2001), 363–384 (electronic).

[WW92] WELCH W., WITKIN A.: Variational surface modeling. *Computer Graphics 26*, 2 (1992), 157–166.

[WW94] WELCH W., WITKIN A.: Free-form shape design using triangulated surfaces. *Computer Graphics Proceedings. Annual Conference Series 1994. SIGGRAPH 94 Conference Proceedings* (1994), 247–256.

[Xu04a] XU G.: Convergence of discrete Laplace-Beltrami operators over surfaces. *Comput. Math. Appl. 48*, 3-4 (2004), 347–360.

[Xu04b] XU G.: Discrete Laplace-Beltrami operators and their convergence. *Comput. Aided Geom. Design 21*, 8 (2004), 767–784.

[YOB02] YAGOU H., OHTAKE Y., BELYAEV A. G.: Mesh smoothing via mean and median filtering applied to face normals. In *GMP* (2002), pp. 124–131.

[Zor] ZORIN D.: Curvature-based energy for simulation and variational modeling. In *Proceedimgs of SMI '05*, pp. 198–206.

[ZT89] ZIENKIEWICZ O. C., TAYLOR R. C.: *The finite element method.* McGraw Hill, 1989. 1.

[ZT97] ZIENKIEWICZ O. C., TAYLOR R. L.: The finite element patch test revisited: a computer test for convergence, validation and error estimates. *Comput. Methods Appl. Mech. Engrg. 149*, 1-4 (1997), 223–254. Symposium on Advances in Computational Mechanics, Vol. 1 (Austin, TX, 1997) 0045-7825.

© The Eurographics Association and Blackwell Publishing 2006.

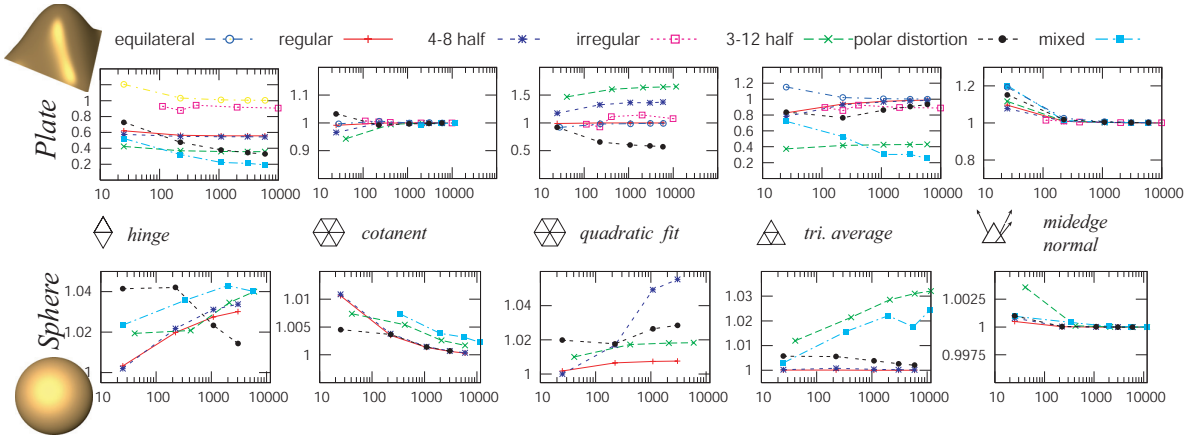

Figure 12: *Linear plate and nonlinear sphere convergence tests for different operators; the vertical axis shows computed displacement normalized by the analytic value.*

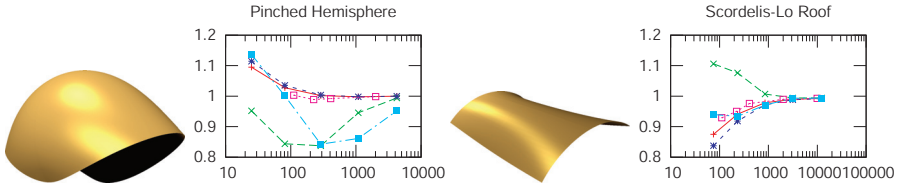

Figure 13: *Engineering obstacle course behavior for our operator.*

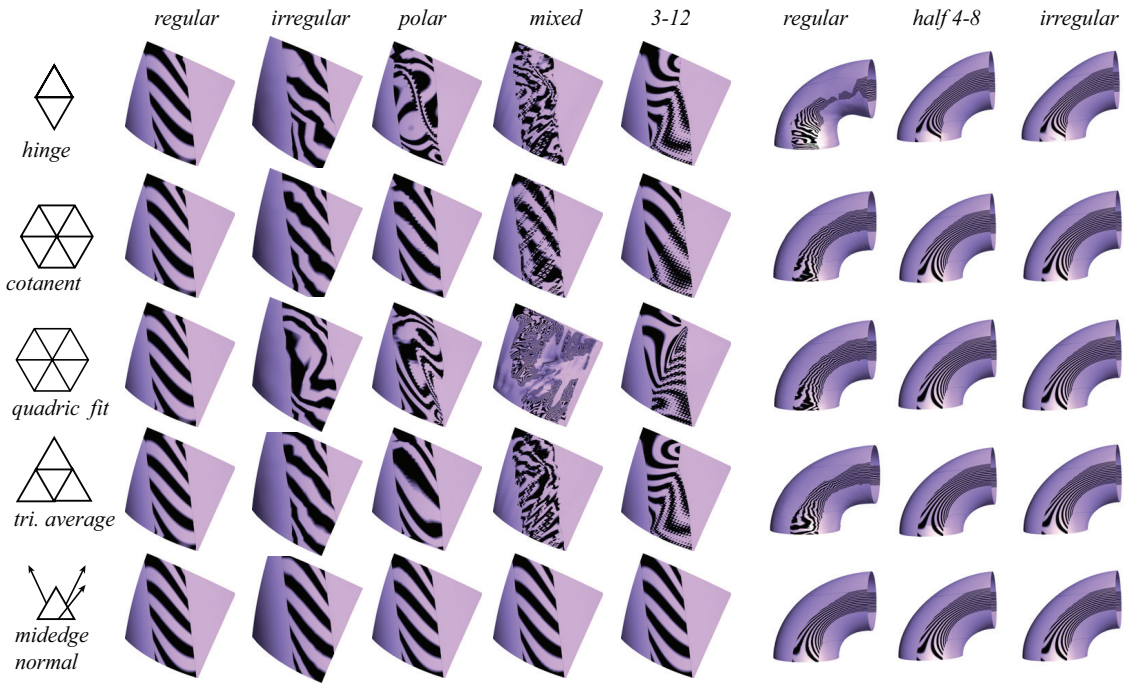

Figure 14: *Left: solutions of the minimization problem* (8) *with boundary conditions sampled from a quadratic cylinder. Right: solution of the same problem for boundary conditions sampled from two circles. In all cases, reflection maps are shown on the surfaces.*

© The Eurographics Association and Blackwell Publishing 2006.