

Energy Minimizers for Curvature-Based Surface Functionals

Pushkar Joshi¹ and Carlo Séquin²

¹ppj@eecs.berkeley.edu, ²sequin@eecs.berkeley.edu
CS Division, University of California, Berkeley

ABSTRACT

We compare curvature-based surface functionals by comparing the aesthetic properties of their minimizers. We introduce an enhancement to the original inline curvature variation functional. This new functional also considers the mixed cross terms of the normal curvature derivative and is a more complete formulation of a curvature variation functional. To give designers an intuitive feel for the preferred shapes attained by these different functionals, we present a catalog of the minimum energy shapes for various symmetrical, unconstrained input surfaces of different genera.

Keywords: surface optimization, bending energy, curvature variation, aesthetic design.




Fig. 1: The control polyhedron of a symmetric genus-2 starting shape (a), with its corresponding minimizers for (b) Bending, (c) MVS and the (d) new MVS_{cross} energy functionals respectively.

1. INTRODUCTION

Freeform surface design is found in CAD systems used for geometric sculpting, consumer products, ship hulls, and automotive shapes. With ever increasing computing power, surface designers will soon be able to use functional optimization as an interactive tool for designing smooth shapes. In a typical aesthetic design task, the designer will provide the input surface along with any geometric constraints such as boundary conditions, position or normal values at some interior positions, or certain parts of space to avoid. All other surface parameters constitute the degrees of freedom that will be adjusted by the optimization system to minimize some “cost” or “energy” functional associated with the surface. The resulting smooth surface will be completely defined by the genus, symmetry, and constraints of the input surface, and the energy functional that is minimized. The exact formulation of the mathematical functional strongly influences the nature of the optimal shapes. Given the same input surface and constraints, different functionals result in different minimizers, each with different aesthetic properties (Fig.1). In this report, we characterize the aesthetic properties of commonly used surface energy functionals. We do so by comparing their respective minimizers for starting shapes of varying genera and types of symmetry.

Traditionally, the functional used most often is the total bending energy. Assuming that the surface is a thin, infinitely stretch-able shell (“Bernoulli’s elastica”), its bending energy is calculated as the area integral of squared principal curvatures. This functional is equivalent to the well-known Willmore energy functional [24]. A popular and publicly available tool for computing the Willmore energy minimizers is Brakke’s Surface Evolver [4], which iteratively refines and optimizes the triangulated input surface to approach the final minimum energy state.

Using the Surface Evolver, Hsu et al. [12] studied and catalogued the Willmore energy minimizing forms for the unconstrained closed input surfaces of genus zero to five. For genus greater than one, the energy minimizers acquire the shape of Lawson surfaces [13] (Fig.1b) or a shape that can be transformed into a Lawson surface by a conformal

Moebius transform[12]. For higher genus (>5), the Lawson surface tends to look more and more like two intersecting spheres, with the intersection line shrinking to a smaller circle area compared to the diameter of the two spheres. Overall, the Willmore energy minimizers tend to be “blobby”; it would be difficult to argue that these forms are the most beautiful shapes that a surface of a given genus could acquire. The extreme differences in curvature found in the Willmore energy minimizers are unacceptable to most designers. Most designers prefer surfaces with a more balanced distribution of curvature – such a surface would have uniform toroidal arms of thickness comparable to that of the rest of the surface features (like the tunnels going through the object).

Out of the above considerations, the Minimum Variation Surface (MVS) functional was born [17]. The authors wanted a functional that penalized curvature change along the lines of principal curvature. They were inspired by the perfect shape of genus zero: the sphere. Arguing that the sphere should thus have a “cost”/“energy”/“penalty” value of zero, they chose to define their energy as any deviation from a sphere. With these premises it was natural to create a functional that calculates the surface integral over the changes in curvature. To keep things simple and efficiently computable (in 1992) the authors only considered in-line curvature changes along the lines of curvature to define their so-called “MVS” energy. This resulted in many pleasing energy minimizing forms [17].

The MVS functional reports a value of zero for cyclides – surfaces whose lines of principal curvature are circles. One drawback of the MVS functional is that *all* cyclides result in a total MVS energy of zero. We introduce and evaluate a new functional that allows us to distinguish between different cyclides and present its energy minimizing forms.

Goal: In a future CAD system, we expect that designers will have a choice among several functionals when designing smooth surfaces. To give designers an intuitive understanding of the aesthetic properties of the different functionals and what kind of surface forms one should expect upon optimization, we present a catalog of energy minimizing shapes (and their associated energy values) for different genus and symmetry constraints.

Since the computation of these shapes turned out to be much trickier than we originally expected, we present the details of our computational approach in Section 3 of this paper. Note that these computations took longer than what one would want in an interactive design system. However, for this report, accuracy and reliability were our foremost concerns, and speed of the optimization system became a secondary issue.

With our benchmark system in place, as future work, we can begin examining approximate methods that will yield similar looking surfaces but at grossly reduced computational costs. Any such system that we or anybody else might develop could then be calibrated and verified against shapes produced in our benchmark system.

2. CURVATURE BASED ENERGY FUNCTIONALS AND RELATED WORK

Most popular surface energy functionals are directly inspired by energy functionals designed for curves. Often, the energy of an infinitesimally small surface area element is defined as the average energy of the surface curves passing through that element. Integrating the energy of the small area element over the entire surface gives the total surface energy. By selecting different curve properties to measure (curvature, change in curvature etc.) and by considering special subsets of the surface curves (geodesics, lines of curvature etc.), we can design various surface energy functionals.

Bending Energy: This well-known functional is a generalization of Bernoulli’s “elastica” energy that measures the square of curvature ($\int \kappa^2 ds$) integrated over the length of a given curve. For curves on a surface, we consider only the component of curvature along the surface normal. This normal curvature ($\kappa_n(\theta)$) is a function of the principal curvatures (κ_1, κ_2) parameterized by the angle θ made with the first principal direction. The average value of the square of normal curvature for *all* geodesic curves passing through a point can be computed by (see [16] for a complete explanation)

$$\frac{1}{\pi} \left(\int_0^\pi \kappa_n(\theta)^2 d\theta \right) = \frac{3}{8} (\kappa_1^2 + \kappa_2^2) + \frac{2}{8} \kappa_1 \kappa_2. \quad (1)$$

For surface optimization, we consider only a subset of the above terms and define the bending energy as an area integral of the surface,

$$\text{Bending Energy} = \int (\kappa_1^2 + \kappa_2^2) dA. \quad (2)$$

There are several equivalent formulations of bending energy, all producing the *same* minimizers. Expanding the term in the above integral gives

$$\int (\kappa_1^2 + \kappa_2^2) dA = \int 4 \left(\frac{\kappa_1 + \kappa_2}{2} \right)^2 dA - \int 2\kappa_1\kappa_2 dA = \int 4H^2 dA - \int 2GdA \quad (3)$$

where H is the mean curvature and G the Gaussian curvature. According to the Gauss-Bonnet theorem, for closed surfaces, the area integral of the Gaussian curvature is a topological constant that depends only on the genus of the surface

$$\int \kappa_1\kappa_2 dA = 4\pi(1 - \text{genus}). \quad (4)$$

Usually, the topological type (that is, the genus) of the surface is unchanged during optimization. Therefore, if we ignore the integral of Gaussian curvature and minimize only the total squared mean curvature, we obtain the same minimizers. Note that bending energy is scale-invariant: uniformly scaling the surface has no effect on its bending energy. Therefore, this functional can be used to optimize closed surfaces without any external constraints. The functional that measures the area integral of the squared mean curvature is the well-known Willmore energy ($\int H^2 dA$) [24]. The stable, local Willmore energy minima have been found for surfaces of different genera [12]. For surfaces of genus 0, the minimal shape is, of course, a sphere, and it has a Willmore energy of 4π and a bending energy of 8π . For genus 1, energy is minimized in the Clifford torus in which the ratio of the two defining radii is equal to $1/\sqrt{2}$ (both, the Willmore and bending energy are 8π). For a higher genus, Kusner [15] has conjectured that the corresponding genus' Lawson surfaces [13] are the global energy minima with a total Willmore energy that lies below a value of 8π .

An interesting insight into the behavior of the bending energy functional is obtained by considering a functional that prefers umbilic points (as those on a sphere) and penalizes surface points where the normal curvature shows a large variation across all directions. Surprisingly, minimizing such a functional is equivalent to minimizing the bending energy (see [22]). That is, if we measure the change of normal curvature across all directions θ ,

$$\frac{1}{\pi} \left(\int_0^\pi \left(\frac{d\kappa_n(\theta)}{d\theta} \right)^2 d\theta \right) = \frac{1}{2} (\kappa_1 - \kappa_2)^2 = \frac{1}{2} (\kappa_1^2 + \kappa_2^2) - (G). \quad (4)$$

This fact becomes useful when we analyze the shape features preferred by the bending energy functional.

For an open surface topologically equivalent to a disc and suspended in a closed 3D boundary curve (with only position constraints at the boundary), this same functional also aims to minimize total surface area. The resulting shapes are the so-called minimal surfaces. The effect of this functional is found in nature: minimal surfaces are naturally found as soap films in equilibrium formed between wire loops. These shapes are characterized by the fact that the mean curvature at every point of the surface is zero.

$$\text{Minimal Surface} \Rightarrow (\kappa_1 + \kappa_2)/2 = 0; \text{ everywhere} \quad (5)$$

For open surfaces with boundary constraints, Hari et al. [11] describe a robust, finite-element-based Willmore energy optimization system. Besides the Surface Evolver, there are several papers that describe a discrete shape operator, usually for producing diffusion flow (specifically mean curvature flow or Willmore flow) for densely triangulated surfaces. While impossible to list every publication concerning discrete operators, we point the reader to some recent work – namely, Bobenko and Schroeder [5] for a novel shape operator that preserves the Möbius invariance of Willmore energy, and Grinspun et al. [10] for a shape operator with strong convergence guarantees. Also see Xu et al. [25] for a discrete operator that studies variation of curvature.

Note that there is a large body of excellent research on interactive surface fairing using linear approximations of the non-linear variational problem. The speed gained is at the cost of dependence on either the input parameterization, or a particular input shape (the so-called “Data Dependent” fairing approaches). Since we are interested in approaches that measure only the intrinsic geometric properties of the surface, we will not list those publications here.

MVS Energy: It has been argued [21] that bending energy may not be the best “beauty functional” for aesthetic surface design. For the surfaces of higher genus, most people prefer a better balance between the diameters of the toroidal handles and the holes between them. This balance is brought about by a more uniform distribution of curvature over the surface. Thus we might obtain a better functional to evaluate the fairness of a curve or surface, if we

try to minimize the integral over the squared *change* of curvature instead. Towards this goal, Moreton and Séquin [17] introduced the “MVS” functional that measures curvature variation by integrating the squares of the derivatives of the principal curvatures in the directions of their respective principal directions. The additional ($\int dA$) is for scale invariance [22].

$$\text{MVS Energy} = \int \left(\frac{d\kappa_1^2}{de_1} + \frac{d\kappa_2^2}{de_2} \right) dA \bullet \int dA. \quad (6)$$

Here e_1 and e_2 are principal curvature directions. The optimal surfaces according to this functional have minimum variation of curvature, hence they are dubbed minimum variation surfaces, thereby giving the functional its name. The MVS functional returns a zero value for surfaces where the principal lines of curvature are exact circles. That is, all the cyclides of various shapes (spheres, cylinders, cones, tori, and even horn tori) are “perfect” surfaces of zero MVS cost. Of course, this is also a drawback of the original MVS functional: *all* cyclides give a zero energy value, and the functional is unable to distinguish between them. Clearly, the asymmetric horn-cyclide should not be judged equally fair and beautiful as a well balanced, rotationally symmetrical torus. In order to obtain such discrimination and rank cyclides according to their aesthetic properties, we need to also take into account the mixed cross terms of the curvature derivative, i.e., the terms dk_1/de_2 and dk_2/de_1 . Towards that goal, we now introduce an enhancement to the original MVS functional.

MVS_{cross} Energy: In addition to considering the change in normal curvature along the in-line direction, we now also consider the change in normal curvature along the orthogonal direction. This leads to the new functional,

$$\text{MVS}_{\text{cross}} \text{ Energy} = \int \left(\frac{d\kappa_1^2}{de_1} + \frac{d\kappa_2^2}{de_2} + \frac{d\kappa_1^2}{de_2} + \frac{d\kappa_2^2}{de_1} \right) dA \bullet \int dA. \quad (7)$$

The reader will notice that the $\text{MVS}_{\text{cross}}$ functional measures the squared 2-norm of the gradient of the principal curvatures with respect to the Riemannian metric (that is, using the principal directions as the basis vectors). Therefore, we believe this functional is a more complete formulation of an energy that measures curvature variation along lines of curvature. In contrast with the original MVS energy, there are only two shapes that have zero $\text{MVS}_{\text{cross}}$ energy: the sphere and an open cylinder (with a circular cross section of any radius). Roughly speaking, the $\text{MVS}_{\text{cross}}$ energy of a surface measures the deviation of the surface from a perfect sphere or a cylinder.

Note that even the $\text{MVS}_{\text{cross}}$ functional measures only a part of the total curvature variation for a surface. A more general curvature variation introduced by [16] measures the variation of curvature for all geodesic surface curves through a point on the surface. An even more general approach is taken by [8] in their definition of six third-order functionals that measure the variation of curvature for *all* surface curves through a point on the surface. A comparison of MVS and $\text{MVS}_{\text{cross}}$ with these more general curvature variation functionals is a topic of future work.

3. EVALUATION FRAMEWORK

Similar to the work of Hsu et al. [12], we push our optimization system to find the true local energy minima for various input shapes. During several stages of our project, seemingly unimportant details of the implementation became crucial factors in limiting the accuracy of our results. Therefore, while our optimization system is conceptually simple, we now describe some important implementation details. In particular, we describe our choices for surface representation and our optimization procedure.

Surface Representation: We use the popular Catmull-Clark subdivision surfaces [6] as a basic representation of the shapes to be optimized. Away from the extraordinary vertices, the surface is simply a collection of bi-cubic B-spline patches connected with G^2 continuity. Therefore, we can use the surface representation directly for Willmore and curvature variation optimization. Also, since Catmull-Clark surfaces have square-integrable curvature everywhere (see [18]), we can use the limit surface directly even near extraordinary vertices for the Willmore energy optimization.

As pointed out by Sabin et al. [19], the curvature values of a Catmull-Clark limit surface diverge near extraordinary vertices. The divergence starts at the patch boundary away from the extraordinary vertex and increases as we get closer to the vertex. In terms of implementation, as we increase the sampling density for a patch adjacent to an extraordinary vertex, the MVS and $\text{MVS}_{\text{cross}}$ energy values will diverge and create numerical instabilities. Therefore, we need to change the limit surface definition for a patch near the extraordinary vertex so as to obtain square-integrable curvature derivatives everywhere on the surface.

We use a method similar to the “flatness” parameter in Biermann et al. [2] and “iron” out the curvature discontinuity near the extraordinary vertex with a flat spot (see Fig. 2). The surface near the extraordinary vertex is defined as a blend between the original Catmull-Clark surface (with the curvature discontinuity) and the limit

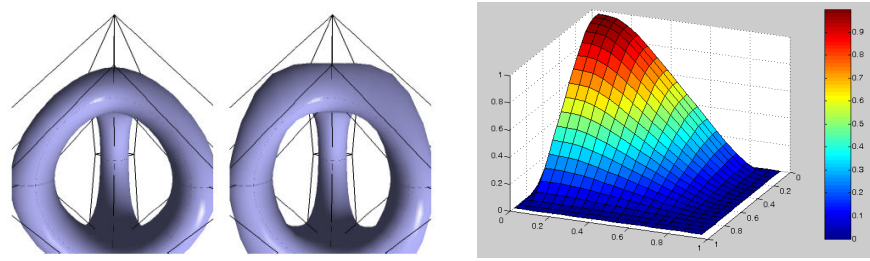


Fig. 2: Original Catmull-Clark limit surface (2a) compared with the modified, blended limit surface (2b) near an extraordinary vertex. The blend function is shown in Fig. 2c. We use the modified surface to obtain the necessary smoothness for curvature variation optimization.

surface projected to the tangent plane of the extraordinary vertex (with no curvature discontinuity). The blend function can be any smooth function that is at least C^3 continuous and has the value 1 (in the limit) at the extraordinary vertex and value 0 (in the limit) at the nearest neighbors. We found that the C^∞ function from Ying and Zorin [26] gave us an initial blend surface with the best aesthetic properties

$$\text{blend}(t) = \frac{e^{(2e^{-1/t})/(t-1)}}{e^{(2e^{-1/(1-t)})/(-t)} + e^{(2e^{-1/t})/(t-1)}}. \quad (8)$$

The control polyhedra of our surfaces have quadrilateral facets, so for a point with parameters (u,v) , the blend function is defined as $\text{blend}(u)*\text{blend}(v)$. The extraordinary vertex has parameters $(0,0)$ – the blend weight achieves value 1 at the extraordinary vertex and value 0 on the patch boundaries not containing the vertex. The resulting blend function is shown in Fig. 2c.

Note that more sophisticated methods like Levin’s scheme [14] or the various manifold-based constructions (e.g.[9], [26]) may produce starting surfaces with aesthetically better shapes without a noticeable flat spot. However, this flat spot is not a major problem for us. As we increase the level of subdivision and move towards a finer control mesh, the size of the flat spot decreases. The additional control vertices introduced with the higher levels of subdivision will be adjusted by the optimization procedure to remove any flatness in the neighborhood of the extraordinary vertex. Thus, our simple smooth surface construction works well for our purposes.

We use the hierarchy inherent to Catmull-Clark subdivision to introduce new degrees of freedom by subdividing the control polyhedron. A denser control polyhedron allows us to approximate non-piecewise polynomial shapes like spheres and tori with piecewise polynomial patches.

We interrogate the limit surface using Stam’s exact evaluation implementation [23]. The Gauss-Legendre quadrature was used to sample quadrilateral patches. The setup of exact evaluation made it possible to take analytic derivatives for energy gradient computations. This code was rather long and complicated for the curvature variation functionals, but ran much faster than the code for an approximated finite-difference based gradient.

Optimization: Our variational problem is solved with a quasi-Newton method that uses energy gradients to approximate a quadratic model of the energy space. The quadratic model indicates an energy minimum and gives a search direction towards that minimum. A line search routine samples the surface energy along that direction to find the first local energy minimum. These two steps are repeated until the energy attains a steady, fixed value. We use the routines in the excellent Toolkit for Advanced Optimization (TAO) [1] for our benchmark system. We did not use methods that required the exact energy Hessian (e.g. Trust Region Methods) for coding convenience.

4. ENERGY MINIMIZING SHAPES

We subjected a set of canonical, symmetric starting shapes to unconstrained optimization (for all three functionals). The differences in the aesthetic properties of the functionals are readily apparent when we compare the minimizing shapes obtained. Our goal was to find the final, optimal shapes and observe the typical shape characteristics of these minimizers. As a test for robustness, we verified that our system was able to reach the same minimum from two different input shapes with the same symmetry, typically one much skinnier and one much fatter than the final minimizer (Fig. 3). We first give a detailed description of our experiences with a torus, and then present a gallery of energy minimizing shapes of high-genus.




Fig. 3: Two drastically different genus-3 input surfaces with tetrahedral symmetry (left, middle) produce the same minimizer (right) upon MVS_{cross} energy optimization.

4.1 Torus Experiments

We performed experiments by sampling an analytic torus surface. The only degree of freedom was the ratio of the two torus radii. These experiments allowed us to compare the optimization of a perfect surface without having to worry about errors brought about by the mesh discretization and piecewise polynomial approximation of a smooth surface. The intuition gained about the behaviour of the functionals was useful in understanding the minimizers for the subsequent, more complicated high genus shapes.

Bending Energy: As expected, all input torus configurations produced the “Clifford Torus” (with a radius ratio of $1/\sqrt{2}$) on Bending energy optimization (see Hsu et al. [12] for a detailed explanation of this shape).

MVS Energy: Since the MVS energy of a torus is zero, all torus configurations produced zero energy.

MVS_{cross} Energy: This energy tries to make the best approximation of a cylinder from a closed torus mesh. Constructing a cylinder from a genus-1 surface without introducing cuts is impossible. In the absence of the scale-invariance term ($\int dA$), an infinitely large torus (of any circular cross section) would be a good piecewise approximation to a cylinder. However, with the scale-invariance term added in, this infinitely large torus gets represented as an infinitely thin torus. Therefore, this infinitely thin torus is the MVS_{cross} minimizer of genus-1 (with energy of approximately 780 – see Fig. 4).

Mixed Energy: Obviously, the genus-1 minimizer of the MVS_{cross} functional is not particularly attractive. However, we now have the ability to combine the Willmore energy with the MVS_{cross} energy to define a new functional

$$\text{Blend Energy} = w_1(\text{Willmore Energy}) + w_2(MVS_{\text{cross}} \text{ Energy}) \cdot \quad (9)$$

The values of w_1 and w_2 can be tweaked so that its genus-1 minimizer is a torus with a desirable ratio of radii spanned by the Clifford torus and an infinitely thin torus. For instance, if $w_1 = w_2$, we get a functional whose optimal torus has a radius ratio of 0.23 (Fig. 5c). If we prefer to obtain a radius ratio of 0.5 (Fig. 5d), we need to set $w_1 = 0.95$ and $w_2 = 0.05$.

4.2 Canonical, Symmetrical High-Genus Surfaces

We present a gallery of minimizers (Table 1) for several high-genus surfaces along with a brief analysis of their shape features. The input surface, which serves only to specify the gross shape and symmetry, is represented by its starting control polyhedron in Table 1.

Bending Energy: As mentioned in Section 2, minimizing the bending energy is equivalent to making every surface point as umbilic as possible (like the surface of a sphere). As a result, the general tendency of the bending energy minimization is to round and “fatten” shapes and make them “bulgier” if possible. In general, holes in the surfaces

shrink and toroidal arms tend to bulge out. In Table 1, we report the total bending energy values found for each minimizer. Given a shape within the proper symmetry class, we obtained the expected Lawson surfaces [13] for all values of the genus; all of the Lawson surfaces found by our system have their Willmore energies less than the conjectured upper bound of 8π [15]. For surfaces with a genus less than 6, these minimal-energy shapes are quite pleasing to look at. However, with increasing genus, these surfaces more closely approximate two spheres intersecting along a circle of alternating tiny pillars and holes, reminiscent of the central portion in Scherk's second minimal surface [20] wrapped into a toroidal ring. In general, there is a sharp variation of curvature values from the round spherical part to the sharply bending intersection area. Most people do not think that this is an aesthetically optimal shape for the higher genus surfaces.




Fig. 4: MVS_{cross} energy values plotted against torus radius ratios. The minimum energy is observed for an infinitely thin torus.




Fig. 5: (a) The Bending energy minimizer (Clifford torus), (b) the MVS_{cross} minimizer (infinitely thin torus), (c) the minimizer of the combined energy with $w1=w2$ (ratio = 0.23), and (d) the minimizer of the combined energy with $w1=0.95$ and $w2 = 0.05$ (ratio=0.5).

MVS Energy: Unlike the case of genus-1 surfaces, higher genus input surfaces produce unique minimizers. Compared to the bending energy minimizers, the MVS energy minimizers have a more uniform distribution of curvature. As a result, these minimizers have thicker toroidal arms that are spaced apart more widely.

MVS_{cross} Energy: The MVS_{cross} functional considers the curvature cross derivative terms and is more discerning than the original MVS. As a result, the overall shape is more uniformly round, and the junctions of the toroidal arms are less “blobby”. For a clear example, see the first two rows of Table 1 (genus 2 and 3), where the toroidal arms of the MVS_{cross} minimizers blend with a bulgier, thicker junction as compared to the MVS minimizers. For the same examples, note that the overall shape of the MVS_{cross} minimizers is rounder than that of the corresponding MVS minimizers.

4.3 Qualitative Discussion of MVS_{cross} Energy

Upon a closer look at the energy values associated with the various MVS_{cross} minimizers, certain characteristics of the functional become evident.

Blended N-way Junctions: Most of the MVS_{cross} energy in a minimizer is due to the blended n -way junctions, where n in our input shapes ranges from 3 to 12. In all these symmetrical shapes, we find that the MVS_{cross} energy increases roughly proportionally with the number of blended junctions. To show this trend, we have calculated the energy cost per blended junction. Tables 2 and 3 list this value for n equal to 3 and 4 for surfaces of various genus. In Table 2, notice that the energy of the blended junction where the toroidal arms come together in a symmetric fashion (top row) is significantly lower than the energy of the blended junction where the toroidal arms come together in an asymmetric fashion (bottom row).

Symmetry: As the value of n in a blended n -way junction goes up, the MVS_{cross} energy seems to increase non-linearly. As an example from Table 1, compare the two minimizers for genus 5, where the more symmetric minimizer with a cube frame (with eight 3-way blends) has significantly lower energy than the minimizer with rotational symmetry about its major axis (with two 6-way blends). This means that the MVS_{cross} energy tends to favor more n -way junctions with lower n , rather than fewer junctions with higher n . This boils down to a rather desirable feature: the MVS_{cross} functional

rewards symmetry. The same effect can be seen for the second and third rows of Table 1, where the more symmetric shape has lower energy than the less symmetric one (favored by the Willmore energy functional). Any surface beauty functional should reward symmetry, and we believe our examples make a good case for the use of the MVS_{cross} functional as an aesthetic energy functional.

4.4 Characterization of Results

In summary, we present a brief, informal, but intuitive description of shape characteristics that are obtained at surface points after optimization with respect to each functional. Clearly, only textual descriptions are inadequate – in order to provide a visual explanation of the characteristics listed below, we presented the shapes of various minimizers in Table 1.

Bending Energy: Every surface point is as umbilic as possible; final shapes can be bulgy with spherical regions.

MVS Energy: *Inline* normal curvature variation is minimized along lines of curvature at every surface point; final shapes have cyclide-like properties (i.e. if possible, curvature lines are circles) with a more uniform distribution of curvature than those obtained by bending energy.

MVS_{cross} Energy: Total normal curvature variation is minimized along lines of curvature at every surface point; final shapes are as spherical or cylindrical as possible and, in general, have rounder envelopes with toroidal arms of more uniform cross-sections than those obtained by MVS energy.

6. DISCUSSION






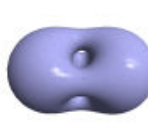







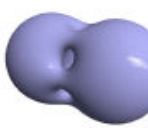


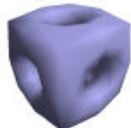
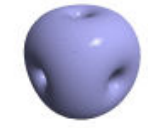
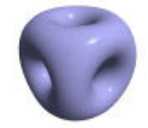
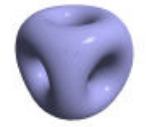





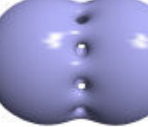

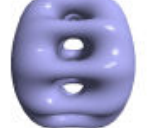
For all the functionals mentioned so far, small local surface perturbations can produce large changes in the surface energy. This sensitivity causes the functionals to be rather stiff or ill-conditioned. The ill-conditioning creates a generic difficulty for optimization of surfaces with dense control meshes. Once every vertex finds itself in a locally optimal neighborhood, the amount of movement it is willing to undergo is very small. As a result, the energy gradient is small, and the optimization of a dense, smooth surface often becomes prohibitively slow. It is possible for the energy gradient to drop to such a small value that the optimization routine prematurely reports convergence to the minimum.

To address this ill-conditioning, we use the multi-resolution aspect of subdivision surfaces. Large-scale shape changes occur at a coarse control mesh resolution and more detailed changes occur at a finer control mesh resolution. Unfortunately, a coarse mesh does not allow us to accurately represent exact surfaces like spheres and cyclides. To a small degree, the discretization errors due to the coarseness of the control mesh influence the shape of the minimum energy shape at that mesh resolution. In theory, this is not a problem: by subdividing the control mesh and introducing more degrees of freedom, we can undo the damage caused by discretization errors. However, in practice, it is not easy to undo the damage: the optimization of the surface with the new, denser control mesh is significantly slower. Therefore, one needs to be careful while selecting the input control mesh resolution – on one hand, the mesh must be coarse enough to bring about fast shape changes, and on the other hand, the mesh must be fine enough to not drive the optimization in a wrong direction.


7. CONCLUSIONS

Computer-aided design tools are gradually becoming more accessible for aesthetic engineering and for artistic shape optimization. We envision the use of functional optimization as a design tool for constructing smooth surfaces. The choice of the particular functional is crucial as it determines the aesthetic properties of the final smooth surface. We have presented a catalog of energy minimizing shapes to provide designers with an intuitive understanding of the effects of the various aesthetic functionals. This catalog will allow designers to predict the types of features they may obtain by performing the time-consuming functional optimization.


In conclusion, we argue for the use of curvature variation as a better aesthetic measure than total curvature. We present a new functional, MVS_{cross} , that is a more complete formulation of curvature variation, thereby resulting in several features that make it attractive as an aesthetic functional.

Genus	Input	Bending	MVS	MVS_{cross}
2				
	Energy:	113($< 40\pi$)	343	2408
3				
	Energy:	142($< 48\pi$)	1548	6416
3				
	Energy:	144	975	4020
5				
	Energy:	195($< 64\pi$)	7331	22688
5				
	Energy:	199	2217	9109
7				
	Energy:	250	3566	18426
11				
	Energy:	351($< 112\pi$)	75720	212031

Tab. 1: Results for high genus shapes: comparison of (left to right) input shape, minimizers of Willmore, MVS and MVS_{cross} energies (with the associated energy values – note these values are in different units). As indicated, the Lawson surfaces we obtained satisfy the theoretical, proposed upper bound for total bending energy [15].



Tab 2: MVS_{cross} energy cost per blended 3-way junction for minimizers with 2, 3, 8, 12 and 12 junctions respectively. The junction energy of the surfaces in the top row with symmetric junctions is lower than that of the surfaces in the bottom row with asymmetric junctions.



Tab. 3: MVS_{cross} energy cost per blended 4-way junction for symmetric minimizers with 2 and 6 junctions respectively.

In our current system, a single shape may take over an hour to fully reach the final optimum. However, most of the substantial shape change occurs in the first few minutes. The remainder is spent in incremental changes that have little visible influence on the shape and only slightly reduce the energy value. For practical design applications, this latter phase could be cut short dramatically. In future work, we plan to explore various methods to make the optimization fast enough for use in an interactive surface design tool.

8. REFERENCES

- [1] Benson, S.; Curfman, L.; More, J.; Sarich, J.: TAO user manual (revision 1.8). Technical Report ANL/MCS-TM-242, Mathematics and Computer Science Division, Argonne National Laboratory (2005) <http://www.mcs.anl.gov/tao>
- [2] Biermann, H.; Levin, A.; Zorin, D.: Piecewise smooth subdivision surfaces with normal control, ACM SIGGRAPH (2000).
- [3] Birkhoff, G. D.: *Aesthetic Measure*. Harvard University Press, Cambridge, Mass. (1933).
- [4] Brakke, K.: The Surface Evolver. *Experimental Mathematics*, 1(2), 1992, 141–165.
- [5] Bobenko, A.; Schroeder, P.: Discrete Willmore flow, *Eurographics Symposium on Geometry Processing*, 2005.
- [6] Catmull, E.; Clark, J.: Recursively generated B-spline surfaces on arbitrary topological surfaces, *Computer-Aided Design*, 10(6), 1978, 350-355.
- [7] DeRose, T.; Kass, M.; Truong, T.: Subdivision Surfaces in Character Animation, *ACM SIGGRAPH*, 1998, 85–94.
- [8] Gravesen, J.; Ungstrup, M.: Constructing Invariant Fairness Measures for Surfaces, *Advances in Computational Mathematics*, 17, 2002, 67-88.
- [9] Grimm, C.; Hughes, J.: Modeling surfaces of arbitrary topology using manifolds, *ACM SIGGRAPH*, 1995.
- [10] Grinspun, E.; Gingold, Y.; Reisman, J.; Zorin, D.: Computing discrete shape operators on general meshes, *Eurographics*, 2006
- [11] Hari, L.; Rivoli, D.; Rubinstein, J.: Computation of open Willmore-type surfaces, *Appl. Numer. Math.*, 37(1-2) 2001
- [12] Hsu, L.; Kusner, R.; Sullivan, J.: Minimizing the squared mean curvature integral for surfaces in space forms, *Experimental Math* 1, 3, 1992, 191–207
- [13] Lawson, H. Jr.: Complete minimal surfaces in S^3 . *Ann. of Math.*, 92(2), 1970
- [14] Levin, A.: Modified subdivision surfaces with continuous curvature, *ACM SIGGRAPH*, 2006
- [15] Kusner, R.: Conformal geometry and complete minimal surfaces, *Bulletin of American Math Society*, 17, 1987
- [16] Mehlum, E.; Tarrou, C.: Invariant fairness measures for surfaces, *Adv. in Comp. Math.*, 8(1), 1998, 49-63
- [17] Moreton, H.; Séquin, C. H.: Functional optimization for fair surface design, *ACM SIGGRAPH*, 1992, 167–176
- [18] Reif, U.; Schroeder, P.: Curvature integrability of subdivision surfaces, *Adv. in Comp. Math.*, 14, 2001
- [19] Sabin, M.; Dodgson, N.; Hassan, M.; Ivrišimtzis, I.: Curvature behaviours at extraordinary points of subdivision surfaces, *Computer-Aided Design*, 35, 2003, 1047-1051, 2003

- [20] Scherk, H. F.: Bemerkungen über die kleinste Fläche innerhalb gegebener Grenzen, *J. Reine Angew. Math. (Crelle's Journal)* 13, 1835, 185-208
- [21] Séquin, C. H.: CAD Tools for Aesthetic Engineering, *Computer-Aided Design & Applications*, 1(1-4), 2004, 301-309
- [22] Séquin, C.; Chang, P.; Moreton, H.: Scale-invariant functionals for smooth curves and surfaces, *Dagstuhl Seminar on Geometric Modelling*, 1993
- [23] Stam, J.: Exact Evaluation of Catmull-Clark Subdivision Surfaces at Arbitrary Parameter Values, *ACM SIGGRAPH*, 1998, 395-404
- [24] Willmore, T.: Note on embedded surfaces, *Anal. Stunt. Ale Univ. Sect. I. a Math.*, 11, 1965
- [25] Xu, G.; Pan, Q.; Bajaj, C.: Discrete surface modelling using partial differential equations, *Computer Aided Geometric Design*, 23, 2006, 125-145
- [26] Ying, L.; Zorin, D.: A simple manifold-based construction of surfaces of arbitrary smoothness, *ACM SIGGRAPH*, 2004

