Minimizers of Curvature-Based Surface Energy

Pushkar Joshi

May 12, 2006

1 Introduction

Significance

Curvature based surface optimization is an important constituent of various geometric modeling and scientific computing tasks. Smooth surfaces are often modelled in a variational setting, where the objective function is used to express the beauty of the shape. Optimization is also used for fitting a smooth, virtual surface to samples of a real, scanned object. Energy minimizing surfaces are also found in the simulation of real-world phenomena such as soap films, cell membranes, liquid interfaces and solder-microchip interfaces. Overall, curvature based surface optimization is a useful technology with a variety of applications.

Functionals

The objective functional most commonly used in surface optimization is the bending energy, also known as the *Willmore* energy. The bending energy measures the total curvature of a surface. Intuitively, this corresponds to the work required to deform a surface from a homogenous flat sheet (which has zero bending energy) to its current shape. The Willmore functional is used to model naturally occurring surfaces such as cell membranes and soap films.

Another functional useful for artistic design is the energy that measures *variation of curvature*. The surfaces that minimize this kind of energy are known as *Minimum Variation Surfaces (MVS)* and the functional is called the MVS functional. One form of the MVS functional was introduced by Moreton and Sequin[15], and is used to model cyclides and cyclide-like smooth shapes. We introduce an enhancement to original MVS functional and compare the new functional to the original.

Difficulties with Surface Optimization

In spite of being an important and useful field, surface optimization has remained a difficult and slow problem. Numerous solutions to specific problems have been introduced, but one has yet to develop a general, efficient and robust tool to optimize an arbitrary input surface into the local minimizer of a variety of functionals. This is understandable, given the difficulties inherent to surface optimization:

- Ill-Conditioning: Optimizing surfaces requires the solution of a non-linear system of equations (∇Energy $= 0$ — analogous to the Euler-Lagrange equations for the given energy). The commonly used energy functionals are very ill-conditioned. That is, small changes in the surface can produce large changes in the surface energy. While this ill-conditioning is useful for quickly smoothing out high-frequency components such as bumps and kinks, it slows down the optimization of low frequency features. That is, once each point on the surface finds itself in a locally optimal neighborhood, the amount of movement it is willing to undergo is very small. As a result, low frequency features take a long time to smooth out, and the surface moves towards the desired optimum shape at prohibitively slow speeds.
- Surface Representation: For optimization, the ideal surface representation would be differentiable, invariant to 'non-shape' parameters (such as parameterization, sampling density, element tesselation quality and aspect ratio), able to denote complex surface topologies (e.g. high genus surfaces with specific symmetry) with only a few degrees of freedom, and very fast to interrogate at every point. Creating such a general yet robust surface representation is often difficult, and the representation is usually tailored to the specific application at hand.
- Complex Energy Landscape: The above two difficulties create a third dif-

ficulty with surface optimization — researchers have a hard time understanding the energy landscape, and often wonder if the energy minimizer they obtained is a local minimum, an inflection point, or the true (and often desired) global minimum.

The difficulties with surface optimization motivate our research. We wish to better understand the energy space of the functionals commonly used in surface optimization. In particular, we wish to study and evaluate the aesthetic properties of different curvature-based energy functionals.

Goal

In this report, we find and analyze the local energy minimizers for a number of canonical input surfaces of varying genus. The results will be driven by accuracy, and not by speed. We build a general, robust and efficient surface optimization system to obtain the energy minimizers. Given these minimizers, we can compare and contrast the different functionals and the 'beauty' of the shapes they produce. This information will eventually be used in a design tool for constructing smooth shapes. The system that we produce will also be useful as a performance and accuracy benchmark for faster optimization systems in future.

2 Related Work

2.1 Willmore Energy

As mentioned above, most of the work in surface optimization centers around minimizing the Willmore energy[22]. Kusner[14] proposed the so-called 'Lawson Surfaces' as minimizers of the Willmore energy for a given genus. Hsu et al.[13] tried to find Lawson Surfaces for a given genus. They used a discrete approximation of Willmore energy (from the Surface Evolver[3]) to produce close-to-optimal surfaces of genus up to five. Hari et al.[12] present a general, robust, finite-element based system for Willmore energy optimization of open surfaces with boundary. Besides the Surface Evolver, various discrete operators for that could be used for bending energy optimization have been described in [20], [23], [6], [10], [4], and [1]. The papers above successfully used discrete operators

for producing duffision flow (specifically 'mean curvature flow'), especially for densely sampled surfaces. Such a diffusion flow is useful for denoising scanned surface data.

2.2 Approximations of Willmore Energy

Parameterization Dependent Functionals For some tasks, it is sufficient to use a parameterization dependent quadratic approximation of the Willmore energy. The resulting linear system can be solved in real-time, which makes this approximation suitable for interactive smooth surface modelers. Unfortunately, this high speed is at the cost of dependence on the particular input parameterization, so different parameterizations of the same shape can produce different optima. These approximations were used in [5], [21], [11]. The parameterization dependent quadratic approximations have also been used in multiresolution variational modelling. For instance, Botsch and Kobbelt[2] describe the basis function for applying the displacement of sparse handles to a dense target surface as the least energy configuration according to the approximate Willmore functional.

Data Dependent Fairing Greiner [8] introduced the a quadratic approximation of the Willmore energy that was independent of the user parameterization, but dependent on a particular surface shape (usually the input shape). This 'data-dependent' energy minimization can be considered a good compromise between the slow, exact Willmore energy optimization and the fast parameterization-dependent energy optimization. Friedel et al. [7] successfully applied the data dependent fairing to subdivision surfaces, using the characteristic map to obtain a smooth enough parameterization near irregular vertices.

Both the quadratic approximations above were acceptable for their main goal, which was to form a smooth surface that followed the coarse user-defined shape, but without unnecessary kinks and bumps. That is, they wanted to smooth out the high frequency features and maintain the low frequency features. On the other hand, we are interested in the gross, global (i.e. low frequency) shape changes to minimize energy. We need independence from any particular user-defined parameterization or even the particular starting shape (other than for defining the symmetry of the desired shape). As a result, we cannot use either of the above quadratic approximations in our system.

2.3 Other Energy Functionals

Moreton and Sequin [15] optimized the variation of curvature (as opposed to the total curvature in the Willmore functional) and introduced the concept of the MVS energy. Higher order functionals (with discrete operators) are used by Xu et al. [23] for producing different types of surface diffusion flow. Parameterization dependent approximations of the MVS energy have also been used in multiresolution variational modeling (e.g. [2]). Greiner also constructed data dependent versions of the MVS energy in [8].

3 System

For clarity and completeness, we describe our choices for optimization routines and surface representation and justify making these choices over others. We remind the reader that the details of these choices are not very important, as our goal is to find the energy minimizing shapes and not the particular path taken to get there.

3.1 Surface Representation

We had three main requirements of the surface representation: **(1)** fast energy and gradient computation, **(2)** conducive to creating smooth shapes with **(3)** efficient storage and few degrees of freedom. We found Catmull-Clark subdivision surfaces to be a good match as a smooth surface representation. We interrogate the limit surface using exact evaluation ([19]): this allows us to treat our surface like any other parametric spline representation, but without the headaches of matching control polygon connectivity. Since curvature of the limit surface is *square integrable* everywhere (even over the irregular vertices — [17]), we can use the limit surface as is for the Willmore energy optimization.

3.1.1 Surface Representation For MVS Optimization

The MVS optimization requires derivatives of curvature, which cannot be reliably computed on subdivision surfaces near irregular vertices. We fix this problem

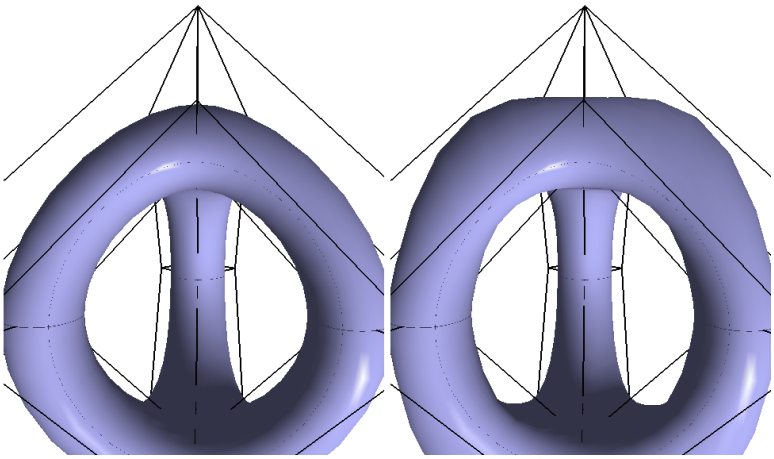

Figure 1: The Catmull-Clark limit surface (left) compared with the fixed limit surface near an irregular vertex (right). Despite the initial 'flat spot', we use the fixed surface representation for MVS optimization.

by modifying the surface representation slightly near the irregular vertices (see Fig. 1). First, we compute the projection of the limit surface near an irregular vertex on that vertex' limit tangent plane. The new surface is then defined as a *smooth blend* between the subdivision limit surface and its projection on the limit tangent plane of the irregular vertex. We use the infinitely smooth (C^∞) blend function ([24]); the blend weight approaches 1.0 (the limit as t approaches zero) at the irregular vertex and zero (the limit as t approaches 1) at all other vertices of adjacent faces:

$$blend(t) = \frac{e^{\frac{2e^{\frac{-1}{t}}}{t-1}}}{e^{\frac{2e^{\frac{-1}{t}}}{t-1}} + e^{\frac{2e^{\frac{-1}{t}}}{t-1}}} \tag{1}$$

While this fix doesn't give a perfect hierarchical surface representation (the same limit surface at different resolutions), it gives us a good enough tensor product parametric surface representation at each subdivision level. This fix was inspired by the fact that for a smooth, dense control polygon, the limit surface in a vertex' local neighborhood is nearly flat. Therefore, the projection of the local limit surface onto the vertex' tangent plane should have negligeble effect, yet allow us to

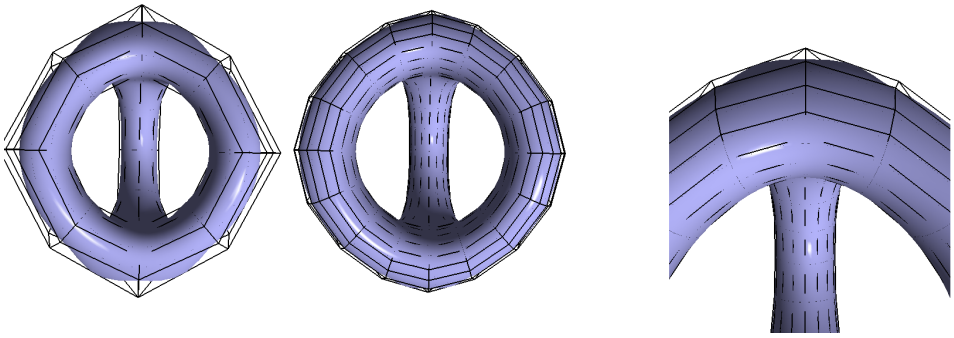

Figure 2: The 'flat spots' near irregular vertices become smaller and unnoticeable with higher resolution (Catmull-Clark subdivision) of the control polygons. On the right is a closeup of the limit surface corresponding to the finer control polygon.

integrate over the entire surface reliably.

We should mention existing work on constructing infinitely smooth surfaces using manifolds ([9], [24]); Ying and Zorin [24] had an elegant construction of C^∞ surfaces using Catmull-Clark surfaces to define the base geometry. The geometry fitting step in [24] requires the calculation of a polynomial surface patch per vertex by a least-squares fit of the limit surface around that vertex. After this step, the connection between the control polygon and the exact surface is broken. Our implementation required an exact relationship between the control polygon and limit surface. Therefore, we did not use the manifolds-based implementation as our smooth surface representation.

The reader will notice that we are creating 'flat spots' near irregular vertices to produce surfaces of sufficient smoothness. At first, this appears counter-intuitive: we are worsening the shape of the surface to get the desired differentiability. However, the shape of the surface improves with the optimization, and as we move towards a finer control polygon, the size of the flat spots also goes down, as shown in Fig. 2.

The reader may also wonder why we used an infinitely smooth blend function as opposed to a simple quartic spline basis function that has necessary parametric smoothness (C^4). As shown in Fig. 3, the quartic spline produces a bigger flat spot than the infinitely smooth function. This fact was further corroborated by

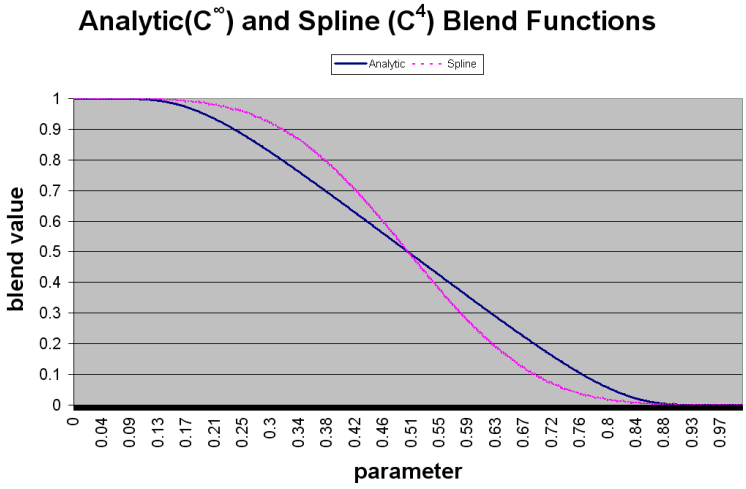

Figure 3: Comparison of C^∞ analytic blend function (solid) with the C^4 spline basis (dotted) — notice how the spline basis forms a bigger 'platform' at the top, which results in a larger, undesirable flat spot than the C^∞ analytic function.

empirical evidence: the energy of the surface obtained using the infinitely smooth blend function was lower than that obtained using a quartic spline blend function. Therefore, we used the C^∞ analytic function as a blend function.

It is important to note that the choice of surface representation does not greatly affect the minimizers produced. Any other smooth surface representation (manifolds, implicit surfaces, points, etc.) could have been used as well. All that our method requires is the ability to compute exact derivatives (that is, with arc-length parameterization) at any surface location. We chose subdivision surfaces because they have a strong research base in general purpose surface modeling tools, and also because of the guarantee of C^2 continuity across all regular patches.

3.2 Hierarchy

We subdivide the control polygon (using Catmull-Clark averaging rules) to introduce more degrees of freedom in the the optimization system. We repeat iterations of optimization (to convergence) followed by subdivision. We perform each such iteration four times — there is no appreciable change in shape between the last

two iterations.

3.3 Optimization

We use the well-known and robust unconstrained non-linear optimization routines in the Toolkit for Advanced Optimization (TAO). In particular, we use the unconstrained quasi-Newton Limited Memory Variable Metric (LMVM) routine as a first pass optimization routine to find a solution. We found LMVM to be faster than non-linear conjugate gradient descent because the number of function evaluations required per iteration was lower on average. We verify the minimum by running a Newton method (with a finite difference Hessian) on the previously obtained solution. We do not compute the exact Hessian for coding convenience; one could also use the simplified, linearized Hessian as in [1].

4 Functionals

Willmore Energy

We use the following expression for the Willmore or thin plate bending energy:

$$\text{Willmore Energy} = \int \left({k_1}^2 + {k_2}^2\right) dA \tag{2}$$

where k_1 and k_2 are the principal curvatures, computed as the eigenvalues of the curvature tensor at the sample location.

There are several variations of this energy. That is, one can define the energy as

$$\begin{aligned} \text{Energy} &= \int \left({k_1}^2 + {k_2}^2\right) dA && (3)\\ &= \int \left(4\left(\frac{k_1+k_2}{2}\right)^2 - 2k_1k_2\right) dA && (4)\\ &= 4\int \left(H^2\right) dA + C_2 && (5) \end{aligned}$$

where H is the mean curvature, and C_2 is a topological constant. The Gauss-Bonnet theorem states that the integral of the Gaussian curvature ($\int k_1 k_2 dA$) is constant for a given genus. Since we are not changing the surface topology during optimization, the use of the constant C_2 is justified. Therefore, Willmore energy minimization reduces to minimizing the square of mean curvature over the surface. For open surfaces with constrained boundaries (i.e. soap films), minimizing surface area is analogous to minimizing the square of mean curvature; this allows one to pose Willmore energy optimization problems as area minimization problems for open, boundary-constrained surfaces ([13]).

MVS Energy

The original MVS functional ([15]) measures the derivative of principal curvatures along their respective principal directions. This is expressed as:

$$\text{MVS Energy} = \int \left(\frac{dk_1}{de_1}^2 + \frac{dk_2}{de_2}^2 \right) dA \cdot \int dA \tag{6}$$

with the additional $\int dA$ providing scale invariance during optimization [16]. Spheres, tori and other surfaces classified as cyclides have zero change of curvature in the respective principal directions; therefore, they have zero MVS energy. Optimizing the MVS energy results in transforming the input shape towards its closest cyclide-like surface. This is desirable as we obtain rounder shapes with thicker, toroidal arms.

However, this can also be a shortcoming of the original MVS formulation: *all* cyclides have zero MVS energy, and the functional cannot distinguish between them. While this fact was known ([18]), the lack of an efficient test bed made evaluating variations of MVS prohibitively slow. Our current faster yet reliable system allows us to aesthetically evaluate different functionals based on variation of curvature. We are particularly interested in an MVS-like functional that is able to distinguish between different cyclides and favors ones that are more symmetrical and have a better balance of positive and negative curvatures. We do this by taking into account the change in principal curvature in the other principal direction. That is, we add the contribution of the $\frac{dk_1}{de_2}^2 + \frac{dk_2}{de_1}^2$ cross terms to the original MVS functional to get:

$$\text{MVS}_{cross}\text{ Energy} = \int \left(\frac{dk_1}{de_1}^2 + \frac{dk_2}{de_2}^2 + \frac{dk_1}{de_2}^2 + \frac{dk_2}{de_1}^2 \right) dA \cdot \int dA \tag{7}$$

We can use this functional to differentiate between different tori, and obtain a new ‘optimum’ torus which has the minor radius about half the major radius (see Table. 1).

It is worth mentioning that our current surface representation (subdivision, i.e. piecewise polynomial surfaces) cannot represent true cyclides: this is because one cannot represent perfect shapes like spheres or tori using B-Splines. However, our system gives us the best B-Spline approximation of cyclides; with additional subdivision, we can add more degrees of freedom for optimization, and the quality of this approximation improves.

5 Results

5.1 Torus Experiments

Please refer to Table 1.

5.1.1 Willmore Energy Optimization

We found that for most torus configurations, the Willmore optimization produced the Clifford Torus as expected (see [13]). However, for a thin torus, we obtained a ‘double sphere’ configuration: a sphere inside another sphere. The energy value was twice the Willmore energy of a sphere, confirming this finding. We are still investigating exactly what range of torus configurations minimize to the Clifford torus, and what range minimizes to this ‘double sphere’.

5.1.2 MVS Energy Optimization

The MVS energy initially approached the Clifford torus (the torus of minimum energy) and then settled on a torus shape close to the Clifford torus. This is understandable: initially, we do not have a perfect torus, so there is a non-zero change of curvature, and therefore MVS Energy. The MVS energy is greater for a coarse mesh and approaches zero with finer torus control polygons. The MVS_{cross} energy produced a torus with the ratio of minor to major radii being about 0.5; we are currently exploring the exact expected configuration of the MVS_{cross} torus.

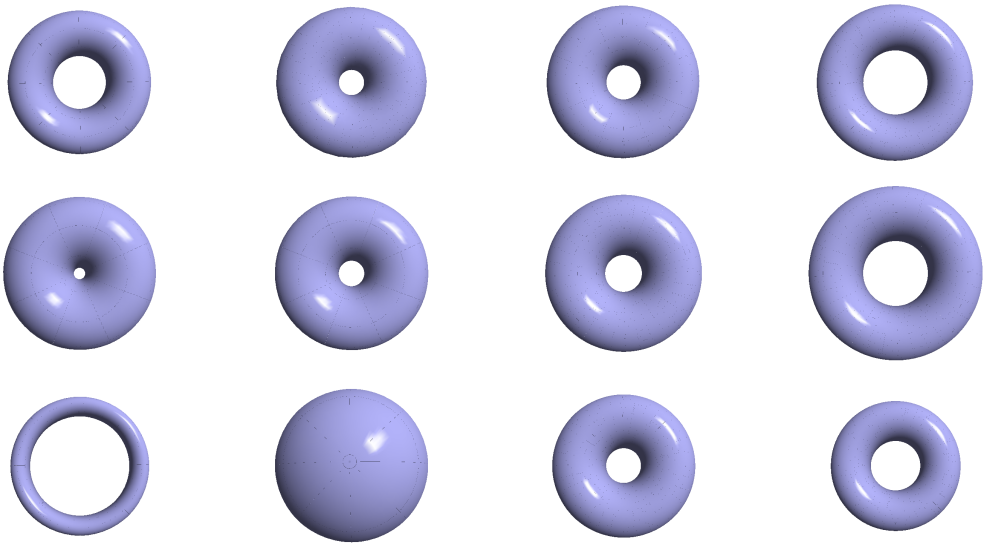

Table 1: Results for various torus configurations: comparison of (left to right) input shape, minimizers of Willmore, MVS and MVS_{cross} energies

5.2 Gallery of Canonical, Symmetrical High-Genus Optimal Surfaces

Please refer to Table 2.

5.2.1 Willmore Energy Optimization

We found that given a shape with proper rotational symmetry, we obtained the expected 'Lawson Surfaces' for any arbitrary genus[14]. However, input surfaces with other types of symmetry (e.g. the genus 3 surface with tetrahedral symmetry, the genus 5 surface with cube symmetry) produced optimal surfaces that maintained the same symmetry.

5.2.2 Willmore Energy Optimization

Like in the case of the torus, the MVS energy optimization first produced shapes close to their corresponding Willmore energy minimizer, and then settled on a

nearby cyclide-like surface. We found a significant difference in shape due to the MVS_{cross} optimization: in particular, we found that the overall shape was rounder, and the toroidal arms were thinner.

6 Conclusion

We are confident we have produced a robust, reliable system for performing curvature-based optimization. We now have a test bed for finding minimizers of various input shapes. Our immediate next goal to use this system to investigate some of the questions raised in the previous section.

While our system is reasonably fast (about 30 minutes per optimization), it is nowhere close to being interactive. In future, we plan on speeding up our optimization by building existing multi-resolution preconditioners (for arbitrary, dense input surfaces) and possibly using robust versions of discrete operators.

References

[1] A. Bobenko and P. Schroeder. Discrete willmore flow. In *Proceedings of Eurographics Symposium on Geometry Processing*, 2005.

[2] M. Botsch and L. Kobbelt. An intuitive framework for real-time freeform modeling. In *Proceedings of SIGGRAPH 2004*, 2004.

[3] K. Brakke. The surface evolver. *Experimental Mathematics*, 1992.

[4] R. Bridson, S. Marino, and R. Fedkiw. Simulation of clothing with folds and wrinkles. In *Symposium on Computer Animation*, 2003.

[5] G. Celniker and D. Gossard. Deformable curve and surface finite-elements for free-form shape design. In *Proceedings of SIGGRAPH'91*, 1991.

[6] D. Cohen-Steiner and J. Morvan. Restricted delaunay triangulations and normal cycle. In *SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry*, 2003.

[7] I. Friedel, P. Mullen, and P. Schroeder. Data-dependent fairing of subdivision surfaces. In *Solid Modeling*, 2003.

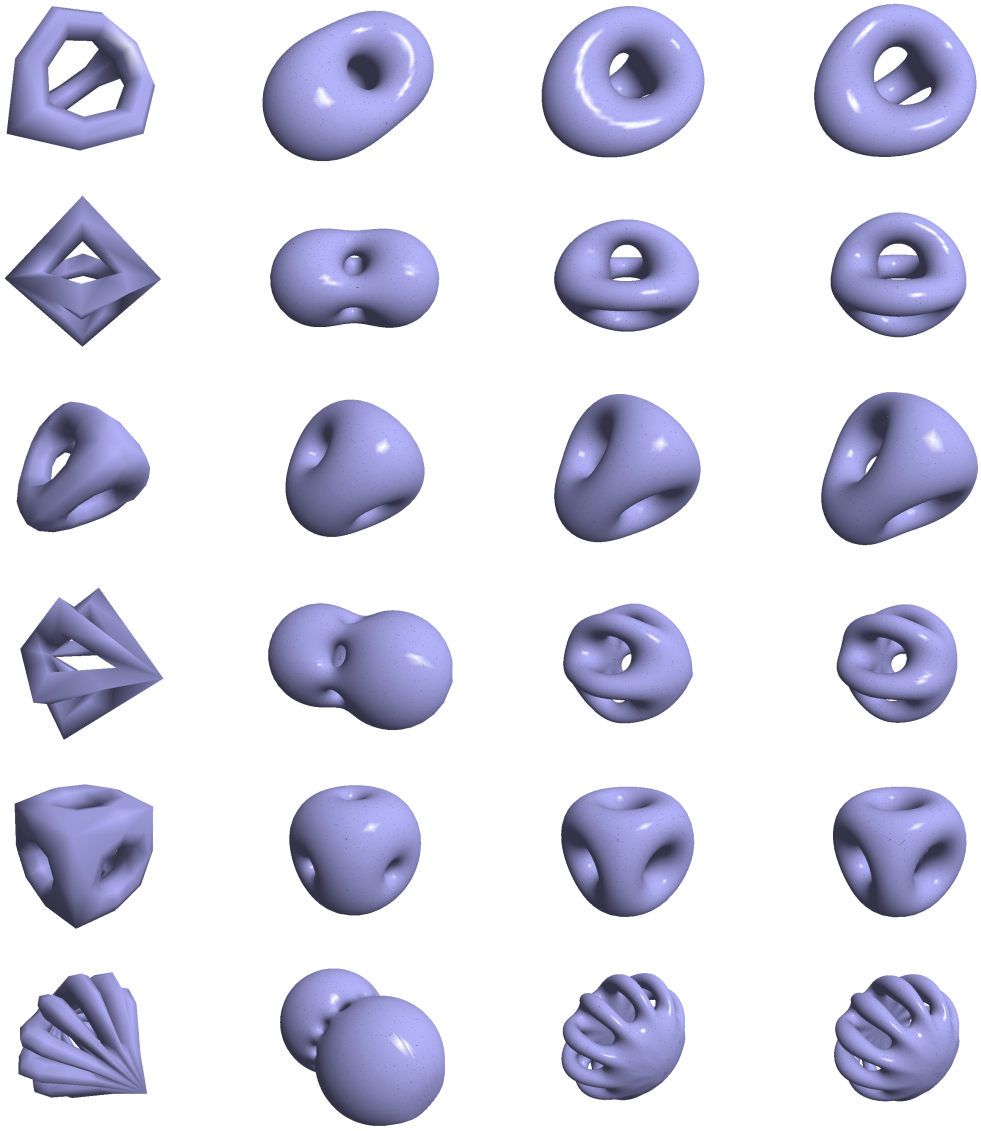

Table 2: Results for various high genus shapes: comparison of (left to right) input shape, minimizers of Willmore, MVS and MVS_{cross} energies

[8] G. Greiner. Variational design and fairing of spline surfaces. *Computer Graphics Forum*, 1994.

[9] C. Grimm and J. Hughes. Modeling surfaces of arbitrary topology using manifolds. In *Proceedings of SIGGRAPH 1995*, 1995.

[10] E. Grinspun, A. Hirani, M. Desbrun, and P. Schroeder. Discrete shells. In *Symposium on Computer Animation*, 2003.

[11] M. Halstead, M. Kass, and T. DeRose. Efficient, fair interpolation using catmull-clark surfaces. In *Proceedings of ACM SIGGRAPH 1993*, 1993.

[12] L. Hari, D., and J. Rubinstein. Computation of open willmore-type surfaces. *Appl. Numer. Math.*, 37(1-2), 2001.

[13] L. Hsu, R. Kusner, and J. Sullivan. Minimizing the squared mean curvature integral for surfaces in space forms. *Experimental Mathematics*, 1992.

[14] R. Kusner. Comformal geometry and complete minimal surfaces. *Bulletin of American Match Society*, 17, 1987.

[15] H. Moreton and C. Sequin. Functional optimization for fair surface design. In *ACM SIGGRAPH*, 1992.

[16] H. Moreton and C. Sequin. Scale-invariant minimum-cost curves: Fair and robust design implements. In *EuroGraphics*, 1993.

[17] U. Reif and P. Schroeder. Curvature integrability of subdivision surfaces. *Advances in Computational Mathematics*, 14, 2001.

[18] C. Sequin, P.Y. Chang, and H. Moreton. Scale-invariant functionals for smooth curves and surfaces. In *Dagstuhl Seminar on Geometric Modelling*, 1993.

[19] J. Stam. Exact evaluation of catmull-clark surfaces at arbitrary parameter values. In *ACM SIGGRAPH*, 1998.

[20] G. Taubin. A signal processing approach to fair surface design. In *Proceedings of SIGGRAPH 95*, 1995.

[21] W. Welch and A. Witkin. Variational surface modeling. In *Proceedings of SIGGRAPH'92*, 1992.

[22] T. Willmore. Note on embedded surfaces. *Anal. Stunt. ale Univ. Sect. I. a Math.*, 11, 1965.

[23] G. Xu, Q. Pan, and C. Bajaj. Discrete surface modelling using partial differential equations. *Computer Aided Geometric Design*, 23, 2005.

[24] L. Ying and D. Zorin. A simple manifold-based construction of surfaces of arbitrary smoothness. In *Proceedings of SIGGRAPH 2004*, 2004.