Analogue of the Total Variation Denoising Model in the Context of Geometry Processing

Matthew Elsey[*]and Selim Esedoḡlu[†]
Department of Mathematics, University of Michigan
Ann Arbor, MI 48109

August 28, 2007

Abstract

We propose a new variational model for surface processing. It is the natural analogue, in the context of geometry processing, of the total variation based image denoising model of Rudin, Osher, and Fatemi. In particular, it admits as solutions surfaces that may have discontinuties in their normals, and can reconstruct shapes with sharp ridges and corners. As the analogue of a perfectly well-posed image model, the proposed surface model differs from those based on ill-posed equations of image processing. Since the proposed geometry regularization is variational, it can be incorporated into a variety of applications, such as surface fairing, 3D object reconstruction from multiple view-points, or shape optimization. Drawing on important previous work of geometers who were interested in extending certain classical theorems from smooth to polyhedral manifolds, we describe in detail the appropriate numerical implementation of the proposed model in the case of triangulated surfaces. We then demonstrate the efficacy of the model with several surface fairing examples.

1   Introduction

Variational and partial differential equations based image denoising models have had great success. The goal of image denoising is to remove noise (in the form of high frequency oscillations) from an image without blurring sharp edges (object boundaries) in the image. These are places where the image

---

intensity function has a discontinuity or very rapid transition between two different values. Often, these models involve degenerate diffusion equations whose diffusivity constant depends on the gradient of the solution.

Recently, some of these models have been extended to the denoising of surfaces [14, 34, 35, 15, 16, 21]. This is a common problem of great interest in computer graphics applications, and is often referred to as surface fairing [27, 39]. It can be a necessary procedure, for example, when real life objects are digitized by a three dimensional scanner by sampling points on the surface of the object. Errors in the coordinates of the points sampled lead to noise in the geometry. The goal of surface denoising is then to remove the oscillations introduced to the surface due to noise. A very important goal in doing so is to preserve the possible sharp discontinuties in the normals of the underlying clear surface, such as creases found on man made objects (e.g. machine parts).

A lot of attention has been devoted to finding the correct analogue of PDE based image denoising models in the context of surface fairing. In particular, there is very interesting previous work on the generalization of the Perona-Malik model [31] to the context of geometry denoising for surfaces [16, 34, 35]. In this paper, we will investigate how to generalize one of the most successful and popular image denoising techniques, namely the total variation based image denoising model of Rudin, Osher, and Fatemi [32] (ROF) to the context of surface fairing. There are several motivations for preferring the ROF model as a candidate for generalization to surfaces over other models such as Perona-Malik. The standard ROF model is a well-posed model with fairly well-understood properties [9]. For example, it has continuous dependence on its data and parameters. Perona-Malik type models, on the other hand, are notorious for having very sensitive dependence on initial data. It seems reasonable to expect that as the analogue of a perfectly well-posed image denoising model, the proposed surface processing model might be better behaved than previously proposed feature preserving surface denoising models that are generalizations of ill-posed image processing techniques.

A second major motivation for considering the ROF model is, as we will explain as a major point of this paper, that it has a natural generalization to surface denoising. The generalizations in the literature about extending e.g. the Perona-Malik type models to surfaces involve making a somewhat arbitrary guess about the exact form of the dependence on the geometry of the surface, especially regarding what combination of principal curvatures should arise. We believe that there is a most natural way to accomplish this goal for the ROF model, making it an appealing option.

2

2   Background and Previous Work

Much attention has been devoted to extending successful variational image denoising models to the context of "geometry denoising", often referred to as "surface fairing". In image processing, the goal of denoising is to remove noise and unnecessary oscillations in the image intensity function, which in the case of grayscale images can be modeled as a function $f(x, y)$ taking its values in the interval $[0, 1]$. The value $f(x, y)$ of the function at a point (x, y) then represents the grayscale intensity of the pixel found at that location. Denoising in these models is accomplished by a diffusion process, described by a parabolic partial differential equation, that can be interpreted as the steepest descent for an energy that penalizes oscillations. The given possibly noisy image is taken to be the initial condition for the diffusion equation; the equation is then solved either for a user-defined duration depending on the amount of noise to be removed, or to stationary state in the presence of fidelity terms.

An overriding concern in image denoising is to preserve sharp "edges" while the noise is diffused away. Edges are the boundaries that separate distinct objects in the scene depicted in the given image. At these locations, the image intensity function $f(x, y)$ has either a sharp transition or a jump discontinuity between the different intensity values of the adjacent objects. Hence, near edges, $f(x, y)$ is expected to have large gradients. To keep edges in the image sharp, the processing PDE are chosen to be nonlinear so that the diffusion constant vanishes, or even becomes negative, at locations where the gradient of the solution has large magnitude. The first example of this type of image processing is given in the revolutionary work of Perona and Malik [31]. The PDE they wrote down is

$$u_t = \nabla \cdot \left(F_\alpha'(|\nabla u|) \frac{\nabla u}{|\nabla u|} \right)$$
$$u(x, 0) = f(x)$$
(1)

where the choice of the function $F_\alpha(\xi) : \mathbb{R} \to \mathbb{R}$ makes all the difference: Perona and Malik's idea is to choose $F_\alpha(\xi)$ to have the following essential properties:

1. $F_\alpha(\xi)$ is even, smooth, and positive for all ξ,

2. $F_\alpha(\xi)$ is convex on $(-k, k)$, and

3. $F_\alpha(\xi)$ is concave on $(-\infty, -k) \cap (k, \infty)$.

Equation (1) arises as gradient descent for the following energy with respect to the L^2 inner product:

$$E(u) = \int F_\alpha(|\nabla u|)\, dx \tag{2}$$

A typical choice for $F_\alpha(\xi)$ is

$$F_\alpha(\xi) = \log\left(1 + \frac{\xi^2}{\alpha^2}\right)$$

The switch in the convexity of $F(\xi)$ that takes place at $|\xi| = \alpha$ corresponds to a switch from forward to backward parabolicity of the PDE (1). Much has been written about the theoretical implications of this feature: It makes (1) ill-posed; there is no satisfactory theory of well-posedness for this PDE. In particular, it is well-known that (1) exhibits extremely sensitive dependence on initial conditions. Nevertheless, the model is extremely effective at its intended purpose of smoothing an image (and thus removing noise and fine detail selectively) without blurring object boundaries, and has become one of the staples of digital image processing.

More recently, PDE based image denoising algorithms have been adapted to the denoising of geometry in the surface fairing application of computer graphics. Noise often gets introduced into the geometry of a surface during the digitization process, as a scanner samples points from the surface of the object and records their 3D coordinates; small jitters and errors during scanning lead to oscillations and noise in the geometry. The goal of surface fairing is to remove such noise from the digitized surface without blurring salient features, such as creases and sharp folds across which the normal to the surface of the original object may be discontinuous; these types of sharp features are very common in, for example, man made objects such as machine parts. Hence, a major concern for surface fairing algorithms is to preserve the relevant discontinuities in the normal field to the surface, and remove noise and spurious detail at the same time.

An important geometric PDE based models for surface denoising has been motion by mean curvature [15]. This well-studied motion turns out to be steepest descent for surface area with respect to the appropriate inner product. It is very effective at removing oscillations from the surface, but makes no attempt at preserving sharp creases or corners. Nevertheless, when sharp features are not present or expected, it is a very effective surface fairing tool, and its proper discrete implementation on triangulated surfaces has led to the interesting work [15, 16] that partially motivates the

present paper. Moreover, it can also be implemented with ease using the level set method [30], and leads to a fairly well understood and numerically manageable second order geometric PDE. Geometric motions with normal velocities given by various other combinations of principal curvatures have also been considered [40, 23, 14, 24].

More recently, previous researchers have looked for ways to generalize edge preserving (i.e. selective) image denoising models, such as the Perona-Malik method, to the task of surface fairing [14, 16, 26]. Since in this context the salient features to be preserved are discontinuties in the normal to the surface, the idea has been that the role of $|\nabla u|$ in (1) and (2) that penalizes oscillations and noise in an image should be played by some combination of curvatures in the context of geometry processing – this is reasonable since curvatures would be expected to be large or singular at rough places on the surface. Which combination of principal curvatures would best serve this purpose has not been addressed clearly in the literature, but several combinations have been tried. For example, references [34, 35] advocate the following energy over surfaces as variational analogues of (1):

$$E(\Sigma) = \int_{\partial\Sigma} F_\alpha \left(k_1^2 + k_2^2 \right) \, d\sigma. \tag{3}$$

In references [34, 35], this energy is shown to lead to a fourth order diffusion equation for the level set function representing the surface $\partial\Sigma$, which is then solved by splitting the fourth order PDE into a system of two second order equations. We point out that such curvature dependent energies arise in other contexts, as well. For example, in image processing, the segmentation with depth problem that concerns itself with reconstructing a 3D scene from a single 2D image of it, has led to models such as [29] by the Nitzberg, Mumford, and Shiota that calls for the minimization of a curvature dependent functional such as (7) over curves in the plane – of course, in the context of curves, the question of what combination of principal curvatures to take does not arise. Another field of image processing that involves curvature dependent functionals is the image interpolation problem known as *inpainting* [25, 5, 10, 19].

It should be pointed out that there are many other approaches to surface fairing. For example, some of the pioneering works in the field, [37, 38, 36], are based on a filtering approach – presumably, they might be connected with a PDE approach such as motion by mean curvature. Another filtering based example is [2]. A more recent example, [20], adapts another very effective image denoising technique, namely the bilateral filtering method, to the surface fairing task. Although these non-variational methods lead to

some of the most impressive denoising results, we believe that variational regularization models are more versatile, since they can be easily incorporated into other reconstruction tasks that appear in, for example, inverse problems.

A related work to ours is that of Alboul and van Damme [1] who use total absolute Gaussian curvature, the surface regularization energy proposed in the present paper, in the different context of connectivity optimization for triangulated surfaces. The fundamental difference between the present paper and [1] (and related works) is that the surface processing we propose maintains connectivity of the vertices but updates continuously their positions (coordinates), leading to a geometric flow of the surface. In [1], on the other hand, vertex positions remain fixed and only the connectivities are updated, leading to a hard combinatorial problem [7]. This type of processing is clearly intended for a very different purpose than that of the present paper.

3 Our Approach

This section is devoted to describing in detail the proposed model for selective surface fairing. First, in Section 3.1, we recall Rudin, Osher, and Fatemi's total variation based image denoising model [32]. Then, in Section 3.2, we provide motivation on how the ROF model should be extended to surfaces. In Section 3.3, we treat the case of "curve denoising" as a stepping stone to our ultimate goal. Section 3.4 finally states our proposed model in light of the insights explained in the previous sections. We also collect in this section some classical theorems from geometry that give substantial support to the proposed model as the natural analogue of the ROF model for geometry denoising.

3.1 Rudin, Osher, and Fatemi Model for Images

A very successful and extremely popular PDE based method in image processing is the total variation based denoising model of Rudin, Osher, and Fatemi. One of its main advantages is its simplicity: It is one of the simplest variational image denoising models that has the all important edge preservation property. Indeed, the model of ROF consists of minimizing the following strictly convex functional:

$$\int |\nabla u| + \lambda \int (u - f)^2 \, dx \tag{4}$$

over all images u. Here, f is the given possibly noisy image, and λ is a user supplied parameter that controls how much noise should be removed: The smaller the value of λ, the larger the amount of signal removed from f. The strict convexity of the model implies, in particular, that the solution exists and is unique. There are a number of interesting minimization procedures for (4) that rely on its convexity, such as steepest descent or duality based methods (numerical solution of the model remains an active research area). In particular, gradient flow for (4) has been studied extensively and given a thorough well posedness theory. It takes the form:

$$ u_t = \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) + 2\lambda(f - u). \tag{5} $$

This flow is to be solved for large times in order to approximate its unique stationary state, which would be the minimizer of (4). Often, this process is substituted by the following regularization procedure:

$$ u_t = \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) $$
$$ u(x, 0) = f(x). \tag{6} $$

This is equivalent to steepest descent for (4) with the choice $\lambda = 0$; of course, the stationary state is then trivial: the constant function. Hence, in practice (6) is run only for a user determined duration of time, and that duration becomes the parameter controlling how much noise is to be removed from the given image. Owing to the convexity of the total variation term that acts as the regularizer, both (4) and (6) have continuous dependence on data f and parameters (λ for (4) and t for (6)). Both procedures lead to very good denoising results and preserve edges in images, as desired.

Advantages of the ROF model over Perona-Malik are two-fold. First, as we already mentioned, both ROF procedures (4) and (6) have complete well-posedness theories and exhibit continuous dependence on their initial conditions, data, and parameters. In practice, this contributes to the robustness of the models. The second major advantage of models (4) and (6) is that they involve a single parameter that needs to be chosen by the user, versus the two parameters (stopping time and α) that need to be chosen for (1).

3.2 Extending ROF to Surface Fairing: Motivation

Compared to previous works, we would like to take a more *global* point of view to the task of generalizing image processing models to the context

of geometry processing. Indeed, wondering what combination of principal curvatures would best serve as the smoothness measure of the surface, in analogy with $|\nabla u|$ in image processing, seems a bit too local. We will instead take our motivations from a different angle.

First, we believe that a useful way to understand image processing models, especially in order to extend them to other contexts such as geometry processing, is to consider the types of images they leave unaltered. In other words, we advocate asking what types images get treated as noise free by these models. In the case of the ROF model in one space dimension, functions with least energy are the monotone functions. Indeed, if we consider functions $f(x)$ defined on the unit interval $[0, 1]$ satisfying given boundary conditions $f(0) = a$ and $f(1) = b$, then as long as the function f is monotone, the regularization term in the ROF model, namely the total variation norm, sees no difference between them: All such functions, whether smooth or discontinuous, are treated the same by the model by assigning them the same energy, namely $|a - b|$. This property is what allows the model to reconstruct discontinuous signals, since it has *no bias* against discontinuities in itself. To summarize, the following properties of the total variation regularization term make it a successful model for image denoising in 1D:

1. $\int_0^1 |f'|\, dx \geq |a - b|$ for every function f with $f(0) = a$ and $f(1) = b$.

2. $\int_0^1 |f'|\, dx = |a - b|$ if and only if f is monotone on $[0, 1]$.

As a warm up towards generalizing the ROF model to surface fairing, we first ask what the reasonable extension of ROF would be to the context of fairing planar shapes, in light of our discussion above. This is the 1D analogue of the surface processing question we set out to address. In this case, the features to be preserved are discontinuities in the unit normal to the curve that forms the boundary of the shape. The regularization term we propose should, of course, still penalize unnecessary oscillations in the direction of the normal (which would be due to the presence of noise), but preserve corners.

We can describe the direction of the normal by considering the angle that the normal makes with the positive x-axis. It is then reasonable to encourage this direction to change in a monotone manner along the curve in order to suppress noise and oscillations in the geometry. In the case of a simple closed curve, the angle is a monotone function of the arc-length parameter if and only if the curve is convex. Another way to state this property of convex shapes, which better lends itself to subsequent generalization, is that *the normal map is one-to-one*. This line of reasoning brings us to the following

8

conclusion: For our denoising purposes, the analogue of monotone functions in the context of planar shape fairing are convex shapes. Thus, the analogue of the ROF model for shape fairing should treat convex shapes as noise-free by assigning to them the least possible energy. The obvious candidate for this purpose is:

$$E_1(\Sigma) = \int_{\partial \Sigma} |k| \, d\sigma \tag{7}$$

where k denotes the curvature of the curve, and $d\sigma$ is the length element. Indeed, (7) is nothing but the total variation of the angle between the tangent to the curve $\partial \Sigma$ and the positive x-axis. It is important to note its connections to the winding number of a closed curve Γ, which is a topological quantity given by

$$\int_{\Gamma} k \, d\sigma = 2\pi N \tag{8}$$

where N is the winding number, an integer. Based on (8), the proposed energy (7) has the following desired properties [17]:

1. $E_1(\Sigma) \geq 2\pi$ for every set $\Sigma \subset \mathbb{R}^2$.

2. $E_1(\Sigma) = 2\pi$ if and only if Σ is convex.

These properties say that all convex shapes are treated as noise-free by model (7), just as we wanted. Moreover, only convex shapes have the least energy, and so have absolutely no noise. In other words, the regularization term E_1 makes no distinction between convex shapes: Whether they are smooth or have corners, they all have the same minimal energy. For instance, a disk, a triangle, and a square all have the same energy. This property endows the model with the ability to reconstruct shapes with discontinuities in the normals.

We now turn to the real focus of this paper: Finding the analogue of total variation denoising for fairing of surfaces. Our starting point is energy (7), but its generalization to surfaces is not completely clear as there are many combinations of principal curvatures that one could imagine taking the place of k. Therefore, we follow the mentality advocated above, and think instead about surfaces that should remain unaltered by the model. In other words, the key to our generalization of (7) will be the requirement that just like (7), our regularization model for surfaces should treat all convex shapes as noise free. A hint comes from (8) that we should look for an analogous topological quantity for surfaces. Such a quantity is of course provided by the Gauss-Bonnet theorem:

Theorem 1 (Gauss-Bonnet) *For a compact, C^2 surface $\partial\Sigma$, we have*

$$\int_{\partial\Sigma} k_G \, d\sigma = 2\pi\chi \tag{9}$$

where k_G is the Gaussian curvature and χ is the Euler characteristic of the surface.

We can now state our proposed surface regularization model. In complete analogy with (8) and (7), based on (9) we propose:

In the context of geometry denoising, the energy

$$E_2(\Sigma) := \int_{\partial\Sigma} |k_G| \, d\sigma \tag{10}$$

is the natural analogue of total variation regularization that appears in the ROF model.

This regularization term has the following desired properties that follow immediately from (9) and the fact that k_G does not change sign (and *the normal map is one-to-one)* for convex surfaces:

1. $E_2(\Sigma) \geq 4\pi$ for every bounded set $\Sigma \subset \mathbb{R}^3$ with C^2 boundary.

2. $E_2(\Sigma) = 4\pi$ if the set Σ is convex.

These two properties imply, as in the case of E_1 for curves, that the energy E_2 makes no distinction between convex surfaces: It does not prefer one to the other based on smoothness, and in particular has no bias against corners and creases. Hence a sphere, a tetrahedron, and a cube all have the same energy, which is minimal among all surfaces. This makes the proposed model capable of reconstructing surfaces with singularities such as ridges and corners that are so common on man-made objects. Furthermore, energy (10) is scale invariant: It remains constant under dilations of the surface $\partial\Sigma$. Hence, the proposed model does not suffer from the "shrinkage" issue that plagues certain other surface processing models, such as motion by mean curvature.

Moreover, we can see that if a surface $\partial\Sigma$ has minimal energy 4π, then in fact Σ has to be convex. Indeed, if Σ is not convex, then there is at least one point $p \in \partial\Sigma$ such that Σ does not lie entirely on a single side of the tangent plane at p. Therefore, there is also another point $q \neq p$ on $\partial\Sigma$ with the same outward unit normal as p, but not contained in the tangent

plane at p. Hence, the Gauss map for $\partial\Sigma$ covers a certain neighborhood on the unit sphere $\mathbb{S}^2$ at least twice. This implies that (10) has to be strictly greater than 4π.

In the subsequent sections, we will discuss denoising models based on the proposed regularization terms of this section, and describe computational techniques for treating them in practice.

3.3 Warm up: Fairing Curves in the Plane

In this preliminary section, we develop the ideas on geometry processing presented above in the simple 1D case of curves in the plane. Although for us this is intended as a warm-up step, it is itself an important problem that has also received considerable attention, since it often arises in computer vision applications, especially in the context of shape recognition. A particularly relevant previous work in this regard is [33]; it develops the analogue of the Mumford-Shah functional [28] for curve denoising and also involves a fourth order PDE. Another is the 2.1D sketch model of Nitzberg, Mumford, and Shiota [29] that involves curvature dependent functionals on curves.

The model we propose, based on the discussions of the previous sections, is the following: Suppose that the curve to be denoised is the boundary $\partial\Omega$ of some region Ω in $\mathbb{R}^2$. Carry out the following minimization:

$$\min_{\Sigma} \int_{\partial\Sigma} |k| \, d\sigma + \lambda\psi(\Sigma, \Omega)$$

The first term in the energy to be minimized above is the regularization term advocated in the previous section. The second term is the fidelity term, which is responsible for keeping the minimizer of the energy close to the original given shape Ω. There are several possibilities suggested in previous studies as a suitable fidelity term, such as:

- $\psi(\Sigma, \Omega) = |\Sigma \triangle \Omega|$, i.e. the area of symmetric difference between the two sets,

- $\psi(\Sigma, \Omega) = \max\{\max_{x\in\Sigma} d_\Omega(x), \max_{x\in\Omega} d_\Sigma(x)\}$, which is the Hausdorff distance between the two sets (where $d_S(x)$ denotes the distance function to the set S).

Noting that the choice of the fidelity term is not the emphasis of this paper, we will simply adopt one of them, say the first one for simplicity, and not dwell on which one is better to use.

The parameter λ sets the relative strength of the denoising. As in image processing applications, there are two options: Either we can set $\lambda = 0$ and carry out gradient descent for the proposed energy (the regularization term) for a certain duration, in which case the stopping time becomes the parameter controlling denoising level, or we can use λ as the parameter that sets the denoising level and find the minimizer of the energy (i.e. carry out gradient descent until stationary state is reached). In either case, we need to be able to implement steepest descent for the regularization term. A very convenient approach to this problem is the level set representation of Osher and Sethian [30]; we show results below with this technique. Another approach is to explicitly represent the boundary of the unknown set as a polygonal curve, and evolve it via its vertices. This second approach has several well-known disadvantages, such as having to keep the vertices more or less uniformly spaced for numerical stability purposes (although this can be done, see e.g. [22]). Nevertheless, we pause below to cover also this method, mainly because it constitutes the 1D analogue of our numerical approach to the 2D surface denoising problem in the next section, where we explicitly represent the surface as a triangulation.

In the level set formulation, we represent the set Σ as the 0-super level set of a function $\phi(x) : \mathbb{R}^2 \rightarrow \mathbb{R}$:

$$\Sigma = \{x : \phi(x) > 0\}$$

so that

$$\partial\Sigma = \{x : \phi(x) = 0\}$$

Then the energy we would like to minimize can be written in terms of ϕ as follows:

$$\int \left| \nabla \cdot \left(\frac{\nabla\phi}{|\nabla\phi|} \right) \right| |\nabla H(\phi)| + \lambda \left(H(\phi) - \mathbf{1}_\Omega(x) \right)^2 \, dx$$

The non-differentiable dependence on ϕ of the integrand makes it necessary to approximate this energy as follows: Let $\Psi_\varepsilon(\xi) := \sqrt{\xi^2 + \varepsilon}$. We minimize instead:

$$\int \Psi_\varepsilon \left(\nabla \cdot \left(\frac{\nabla\phi}{|\nabla\phi|} \right) \right) |\nabla H(\phi)| + \lambda \left(H(\phi) - \mathbf{1}_\Omega(x) \right)^2 \, dx$$

Euler-Lagrange equations for similar energies involving regularization terms of this type can be found, for example, in [41]. It turns out they lead to the following gradient descent equations:

$$\partial_t\phi = |\nabla\phi|\nabla \cdot \left(\frac{\nabla\phi}{|\nabla\phi|} \Psi_\varepsilon(k) - \frac{1}{|\nabla\phi|}(Id - P_{\frac{\nabla\phi}{|\nabla\phi|}}) \left(\nabla \left(\Psi'(k)|\nabla\phi| \right) \right) \right)$$

where we used the following notation:

$$k := \nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \quad \text{and} \quad P_v w = (v \cdot w) v$$

the first being the curvature of the level sets, and the second is the projection operator (onto a vector). Figure 1 shows an example computation with this level set implementation.
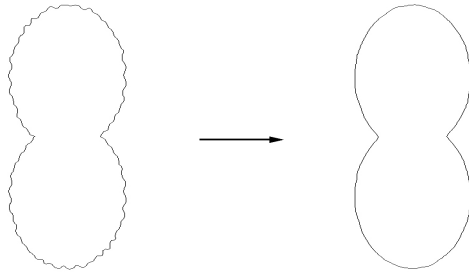

Figure 1: Steepest descent for (7) using the level set method. Oscillations on the curve are diffused out, but the kinks in the middle are fairly well preserved.

We now turn to the alternative method of representing the boundary curve $\partial \Sigma$ of the set Σ explicitly as a polygon in the plane. Let the vertices of the curve be given as

$$\big\{ (x_0, y_0), (x_1, y_1), \ldots, (x_n, y_n) \big\}$$

listed in counter-clockwise direction, so that (x_j, y_j) is connected to (x_{j+1}, y_{j+1}) with an edge of the polygon, and $(x_0, y_0) = (x_n, y_n)$ so that the curve is closed. Then, the energy we minimize is the following discrete analogue of (7):

$$\begin{aligned} E(x_0, \ldots, x_n; y_0, \ldots, y_n) &= \sum_{i=0}^{n-1} |\theta_i| \\ &= \sum_{i=0}^{n-1} \left| \cos^{-1} \left(\frac{e_{i-1} \cdot e_i}{|e_{i-1}| \, |e_i|} \right) \right| \end{aligned} \tag{11}$$

where e_i are the edges of the polygon:

$$e_i = (x_{i+1} - x_i, y_{i+1} - y_i)$$

13

pointing in the counter-clockwise direction. The e_i are approximations to the tangents to the curve represented by the polygon, and θ_i are the changes in the tangent angle; see Figure 2 for an illustration. The curvature of the polygonal curve is concentrated at its vertices; the integral of curvature in a small neighborhood of vertex (x_i, y_i) is given by θ_i. Naturally, formula (11) takes curvature to be 0 along the edges. A very useful and important property of approximation (11) is the following:

1. The energy (11) of any closed polygonal curve is at least 2π, and

2. A polygonal curve has least energy, namely 2π, if and only if it is the boundary of a convex shape.

Thus, we see that not only the continuum model (7) but also its discretization given by the discrete energy (11) satisfies our requirements for a good curve denoising model. In other words, the discrete version of the energy satisfies the important properties of the continuum model *exactly*, and not merely in some limiting sense. We will see in the next section that previous work by geometers on the subject allows the same to be done for the triangulated version of our proposed surface denoising model.

It is now easy to take variations of energy (11) with respect to the coordinates (x_i, y_i) of the vertices v_i forming the polygonal curve. An important point worth mentioning is the non-smoothness of energy (11) whenever $\theta_i = 0$ for some i. This problem can be addressed just as in the standard ROF model by regularizing energy (11) as follows:

$$E = \sum_{i=0}^{n-1} \sqrt{\theta_i^2 + \varepsilon} = \sum_{i=0}^{n-1} \sqrt{\cos^{-1}(\xi_i)^2 + \varepsilon} \tag{12}$$

where we define ξ_i as

$$\xi_i = \frac{e_{i-1} \cdot e_i}{|e_{i-1}| \, |e_i|}.$$

Taking variations of (12) with respect to x_i gives

$$\frac{\partial E}{\partial x_i} = \sum_{j=i-1}^{i+1} \frac{\partial(\sqrt{\theta_j^2 + \varepsilon})}{\partial x_i}$$

$$= \sum_{j=i-1}^{i+1} \frac{\cos^{-1}(\xi_j)}{\sqrt{\cos^{-1}(\xi_j) + \varepsilon}} \cdot \frac{-1}{\sqrt{1 - \xi_j^2}} \cdot \frac{\partial \xi_j}{\partial x_i}$$

14

where the terms $\frac{\partial \xi_j}{\partial x_i}$ are given by

$$\frac{\partial \xi_{i-1}}{\partial x_i} = \frac{1}{|e_{i-2}||e_{i-1}|} \left[x_{i-1} - x_{i-2} - \xi_{i-1} \frac{|e_{i-2}|}{|e_{i-1}|} (x_i - x_{i-1}) \right],$$

$$\frac{\partial \xi_i}{\partial x_i} = \frac{1}{|e_{i-1}||e_i|} \Big[x_{i+1} - 2x_i + x_{i-1}$$

$$- \xi_i \left(\frac{|e_i|}{|e_{i-1}|} (x_i - x_{i-1}) + \frac{|e_{i-1}|}{|e_i|} (x_i - x_{i+1}) \right) \Big],$$

$$\frac{\partial \xi_{i+1}}{\partial x_i} = \frac{1}{|e_i||e_{i+1}|} \left[x_{i+1} - x_{i+2} - \xi_{i+1} \left(\frac{|e_{i+1}|}{|e_i|} (x_i - x_{i+1}) \right) \right].$$

$$(13)$$

The expression for $\frac{\partial E}{\partial y_i}$ is exactly the same as the above formulas for $\frac{\partial E}{\partial x_i}$, only with x's replaced by y's. Define n_i, an approximation to the outer unit normal associated with the i-th vertex, as

$$n_i = \frac{\big(-(y_{i+1} - y_{i-1}), (x_{i+1} - x_{i-1}) \big)}{\big| (x_{i+1} - x_{i-1}, y_{i+1} - y_{i-1}) \big|}.$$

Gradient descent for (12) can now be written as

$$\frac{d}{dt}\big(x_i(t), y_i(t) \big) = \frac{1}{L_i} \left\{ \left(\frac{\partial E}{\partial x_i}\big(x_i(t), y_i(t) \big), \frac{\partial E}{\partial y_i}\big(x_i(t), y_i(t) \big) \right) \cdot n_i(t) \right\} n_i(t).$$

where L_i is the length element associated with vertex v_i, e.g. $L_i = \frac{1}{2}(|e_{i-1}| + |e_i|)$. A sample calculation is shown in Figure 3. We note that projection of the flow onto an estimate of the normal direction is presumably unnecessary since the gradient of an energy that depends only on shape should already point in the normal direction.
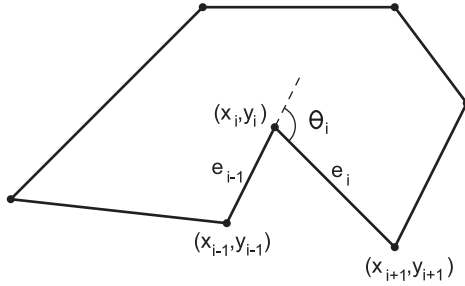

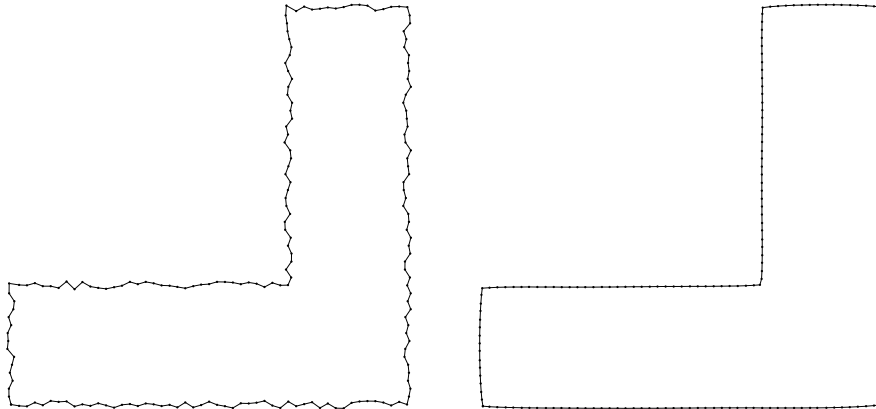Figure 2: Polygonal version of the curve denoising model.

Figure 3: Polygonal version of the curve denoising model in action. The noisy L-shaped curve is denoised, but the result maintains the relevant sharp corners.

3.4 The Proposed Model: Fairing Surfaces

In this section we focus on the main point of the present paper: The analogue (10) of the ROF model in the context of geometry processing that was introduced in Section 3.2. We will record some important theorems from geometry that indicate how the proposed model is appropriate for its intended purpose. We will also discuss the numerical implementation of the model in detail, which relies on some important previous work as well.

Our proposed regularization term (10) appears in previous geometry literature, and is often called, unsurprisingly, *total absolute Gaussian curvature*. Furthermore, geometers in fact studied surfaces with minimal total absolute Gaussian curvature (within a given topological class) previously in detail; these surfaces are said to be *tight*. A theorem, very important for our purposes, establishes a connection of tight surfaces, and hence energy (10), with a certain generalized notion of convexity. This notion is referred to as the *two-piece property*. It is defined as follows. A surface is said to have the two-piece property if any plane divides the surface into at most two connected components. Boundaries of convex regions clearly satisfy this property; however, they are not the only ones. The two-piece property is a generalization of the notion of convexity because there are surfaces with different topology that satisfy it, such as certain tori. These surfaces are, in some sense, "as convex as possible" given their topological type. In particular, there are no spurious oscillations on them. It is intuitively reasonable that the two-piece property precludes spurious oscillations, or noise, on sur-
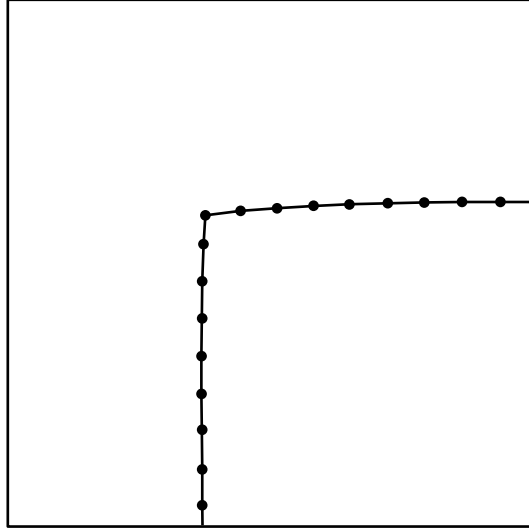
Figure 4: Zoom in onto one of the corners of the denoised curve from Figure 3.

faces: If the surface is rough with bumps, one can take a plane that just grazes the bumps and splits the surface into many connected components, not just two. The following remarkable theorem connects this intuitive notion back to tight surfaces and our proposed regularization (10):

Theorem 2 (from Banchoff & Kühnel [4]) *A compact surface $\partial\Sigma$ is tight if and only if it has the two-piece property.*

This theorem provides further strong support for our proposed model of surface processing. In particular, it elucidates other surfaces that are treated as noise-free by the regularization term (10). Note that we would not expect steepest descent for (10) to be able to take us from, say, a tight torus to a convex shape, even though a convex shape would have energy 4π versus 8π for the torus. This is because we would expect steepest descent for (10) to generate a continuous evolution with respect to some reasonable metric, and a tight torus is presumably a local minimizer in that metric. From an applications point of view, this is entirely acceptable.

The main task before us now is the numerical implementation of steepest descent for energy (10). The first step is to choose a representation for the surface. As we have seen in the case of curves, we have two options: An implicit representation by the level set method, or an explicit representation

given by a triangulation. We found that although it is very easy to express energy (10) in terms of a level set function, its variation (Euler-Lagrange equation) turns out to be quite complicated and a difficult expression to work with. Indeed, the expression for (10) in terms of a level set function ϕ representing the surface $\partial\Sigma$ as $\partial\Sigma = \{x : \phi(x) = 0\}$ can be obtained by putting together two terms that appear in previous level set based models of surface denoising. Namely, in [34] the following level set representation of *total curvature* $k_T^2 = k_1^2 + k_2^2$ is used in the energy density in their proposed model of surface denoising:

$$k_T^2 = \sum_{i,j} \partial_j \left(\frac{\partial_i \phi}{|\nabla \phi|} \right) \left\{ \partial_j \left(\frac{\partial_i \phi}{|\nabla \phi|} \right) - \frac{\partial_j \phi_j}{|\nabla \phi|^2} \right\} \tag{14}$$

Moreover, in [18] the following level set representation of *Willmore energy density* $k_M^2 = (k_1 + k_2)^2$, which also appears in surface processing models [13, 6], is derived:

$$k_M^2 = \left[\nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \right]^2 . \tag{15}$$

The difference of (14) and (15) yields Gaussian curvature k_G:

$$k_G = \frac{1}{2}(k_M^2 - k_T^2) \tag{16}$$

This expression can be somewhat simplified to become:

$$k_G = \frac{1}{|\nabla \phi|} \sum_{i,j} (\partial_i \phi) (\partial_j \phi) \det \left(\mathrm{Hess}(\phi)_{i,j} \right) \tag{17}$$

where $\mathrm{Hess}(\phi)_{i,j}$ denotes the (i,j)-th cofactor submatrix of the Hessian of the function $\phi : \mathbb{R}^3 \to \mathbb{R}^3$. See [12] for another derivation. Note that surface processing flows with integrals of (14) or (15) are equivalent as long as the surface maintains its topology, since the integral of their difference is a topological quantity.

Despite the ease with which k_G is obtained in terms of ϕ above, and although the resulting expression can be somewhat simplified as shown in formula (17), going further and taking the variation of (10) with respect to ϕ yields many complicated terms. Therefore, we will proceed in this paper with triangulated surfaces, and leave the very interesting project of a level set based implementation to a subsequent work (for example, the splitting procedure used in [34, 35] might be worth exploring). There are at least four

benefits to using triangulated surfaces. First, they arise very often in computer graphics applications. Second, there is a very developed theory for our regularization term (10) in the context of triangulated surfaces, which was developed by geometers interested in extending some classical results from geometry of smooth surfaces to surfaces with singularities. Third, compared with level set methods, working with triangulated surfaces sometimes leads to faster computations. And finally, in the denoising applications considered in the present paper, one would not expect (or desire) the evolution to take the surface far enough from its initial state to involve a topological change; therefore, one of the main advantages of the level set method (graceful handling of topological changes) is not so relevant. On the other hand, an important disadvantage of working with triangulations is the need to keep vertices more or less uniformly spaced and the triangles fairly regular on the surface, for numerical stability purposes.

As we already mentioned, several works by previous authors enable us to work with regularization (10) very effectively in the context of triangulated surfaces. First of all, the important work [3] shows how to define the Gaussian curvature of a triangulated surface so that the Gauss-Bonnet theorem, which was important for the motivation leading to our proposed regularization, holds exactly. It would be reasonable for Gaussian curvature to vanish on the faces and along the edges of a triangulation, since it would be for one of the principal curvatures at such locations. The only points left are the vertices, where the curvature should be concentrated, as delta functions. Indeed, the formula in [3] assigns curvature to the vertices, and leads to the following very simple discrete version of $\int_{\partial \Sigma} k_G \, d\sigma$:

$$\sum_i (2\pi - \Theta_i) \tag{18}$$

where Θ_i is the *total angle* at the i-th vertex v_i of the triangulation, defined as follows:

$$\Theta_i = \sum_{f_j \text{ s.t. } v_i \in f_j} \theta(v_i, f_j)$$

where $\theta(v_i, f_j)$ denotes the angle formed at vertex v_i inside the face (triangle) f_j that it belongs to. We will use the notation $k_{G,i} := 2\pi - \Theta_i$; this quantity represents the integral of Gaussian curvature in a neighborhood of the vertex v_i. See Figure 5 for an illustration.
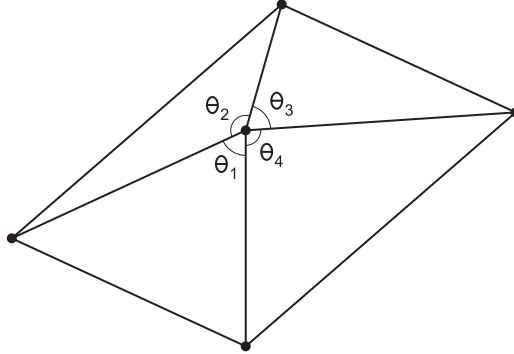
In summary, the theorem in [3] is the following:

Figure 5: Total angle $\Theta = \theta_1 + \theta_2 + \theta_3 + \theta_4$ formed at a vertex of the triangulation.

Theorem 3 (Banchoff [3]) *The following discrete version of Gauss-Bonnet theorem holds for closed triangulated surfaces:*

$$\sum_i k_{G,i} = \sum_i (2\pi - \Theta_i) = 2\pi\chi$$

where χ is the Euler characteristic of the surface.

At this point, a reasonable guess would be that the discrete analogue for our proposed regularization term (10) is obtained simply by replacing the summands in (18) with their absolute values:

$$\sum_i |k_{G,i}| = \sum_i |2\pi - \Theta_i| \qquad (19)$$

One can then express the angles $\theta(v_i, f_j)$, and then the total angles Θ_i, and finally the entire discrete energy (19) in terms of the coordinates (x_i, y_i, z_i) of the vertices v_i of the triangulation, obtaining a system of ODEs that represent gradient descent for the discrete energy (19). This was indeed our first attempt at implementing the proposed denoising model. However, this idea fails completely. In many examples, the steepest descent terminates too soon, reaching a stationary state that is clearly not optimal, leaving intact small irregularities on the surface that should have been removed. The underlying cause of all these problems turns out to be the unfortunate fact that the discrete energy (19) admits minimizers that violate the implications of classical theorems that apply to the continuum version (10) of the energy. In other words, important facts such as Theorem 2 concerning the continuum

20

energy (10) do not extend to this specific discretization (19). In particular, surfaces with minimal discrete energy (19) need not have the two-piece property. For example, Figure 6 shows a triangulated surface, obtained by depressing one face of a cube, that turns out to have minimal discrete energy (19) of 4π; yet this surface clearly lacks the two-piece property.
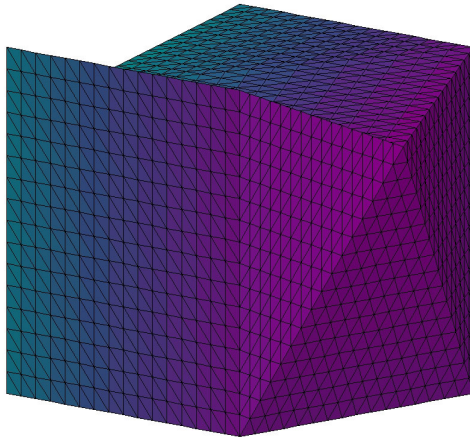

Figure 6: A surface with minimal discrete energy (19) of 4π that clearly lacks the two-piece property. This is because discretization (19) of the proposed regularization term (10) turns out to be inappropriate, as important theorems such as Theorem 2 do not extend to it. Consequently, in applications, gradient flow leaves many undesirable artifacts on a noisy surface. The situation is rectified by using the discrete version of (10) given by Banchoff and Kühnel in [4] for which the discrete analogues of the important theorems for (10) do hold.

Delving further into previous geometry literature concerning tight surfaces, we see that in [4] the *correct* discrete version of energy (10) is given. It consists of defining at every vertex v_i the *positive curvature* (weighted by an area element) $k_{G,i}^+$ as follows. Let $\text{Star}(v_i)$ denote the set of all vertices of the triangulated surface connected to v_i by edges. Define $k_{G,i}^+$ as follows:

1. If v_i lies on the surface of the convex hull of $\text{Star}(v_i)$, let

$$k_{G,i}^+ = 2\pi - \Phi_i$$

where Φ_i is the total angle at v_i on the triangulated surface that forms

21

the boundary of the convex hull of $\text{Star}(v_i)$.

2. If v_i lies in the interior of the convex hull of $\text{Star}(v_i)$, let $k_{G,i}^+ = 0$, i.e. in this case we take $\Phi_i = 2\pi$.

Then, one defines the *negative curvature $k_{G,i}^-$* at vertex v_i as $k_{G,i}^- = k_{G,i}^+ - k_{G,i}$, and finally, the *absolute curvature $|\kappa|_i$* at v_i as $|\kappa|_i := k_{G,i}^+ + k_{G,i}^-$. Then, the correct version of (10) for triangulated surfaces, i.e. its proper discretization given by Banchoff and Kühnel, turns out to be the following:

$$
\boxed{
\begin{aligned}
E_2(v_1,\ldots,v_n) &:= \sum_i |\kappa|_i = \sum_i \left(k_{G,i}^+ + k_{G,i}^- \right) \\
&= \sum_i \left(2\pi - 2\Phi_i + \Theta_i \right)
\end{aligned}
}
\tag{20}
$$

The following theorem from [4], concerning the "discretization" (20) of (10) on triangulated surfaces, is very important for our application:

Theorem 4 (from Banchoff & Kühnel [4]) *A closed triangulated surface with Euler characteristic χ has minimal discrete energy (20) of $2\pi(4-\chi)$ among surfaces topologically equivalent to itself if and only if it has the two-piece property.*

This theorem assures us that if the discrete version of our proposed denoising model (10) is implemented according to (20), then the properties of the model explained in Section 3.2 that originally motivated us will indeed hold. In particular, the two-piece property will ensure the absence of spurious features that were left on the surface with the unsatisfactory implementation (19).

Our proposed denoising model calls for expressing energy (20) in terms of the coordinates of the vertices $v_1,\ldots,v_n$, and carrying out gradient descent by updating the surface via the v_i. As far as we know, the geometric motion that results from gradient descent for energy (10) had not been considered previously, and constitutes a new model for surface fairing. Even with all the powerful previous work explained above that we can draw upon in the context of triangulated surfaces, implementing this geometric motion remains challenging. Indeed, although all the quantities appearing in (20) are straight-forward to express in terms of the coordinates of the vertices v_i, the definition of $k_{G,i}^+$ is piecemeal and involves a switch, depending on the convex hull of the immediate neighborhood of v_i, whose nature might of

course change during the evolution. This might lead to a non-smooth optimization problem in terms of the coordinates. Such is the case, for example, with the ROF model in image processing: The energy is non-differentiable, so that its numerical treatment has been the topic of much research.

In this paper, we take a first step towards implementing model (10): We carry out gradient descent for discrete energy (20) by ignoring all possible non-differentiability issues for the time being. This may lead to oscillations about the singular points, but we see no evidence of this in our computations. A more theoretically justified minimization technique will form the basis of further work. In particular, the duality based methods [11, 8] originally developed for the ROF model which allow avoiding the non-differentiability issues, might have analogues in this setting (although we do not expect convexity of the energy).

We now describe our straight forward implementation. For each i and j such that $v_i \in f_j$, let $\ell_1(i,j)$ and $\ell_2(i,j)$ be the corresponding vertex indices such that $\{v_i, v_{\ell_1(i,j)}, v_{\ell_2(i,j)}\}$ constitute the vertices of face f_j of the triangulated surface, listed in order of positive orientation. For each i, let $\{\tilde{f}_{i,1}, \ldots, \tilde{f}_{i,m(i)}\}$ denote the faces of the triangulation that forms the surface of the convex hull of $\mathrm{Star}(v_i)$ whose vertices are a subset of the original set of vertices $\{v_i\}$ of the given triangulated surface. And for each i and $j \in \{1, \ldots, m(i)\}$, let $\tilde{\ell}_1(i,j)$ and $\tilde{\ell}_2(i,j)$ denote the vertex indices such that $\{v_i, v_{\tilde{\ell}_1(i,j)}, v_{\tilde{\ell}_2(i,j)}\}$ constitute the vertices of face $\tilde{f}_{i,j}$, listed in order of positive orientation. Then, energy (20) can be expressed as

$$E_2(v_1, \ldots, v_n) = 2\pi n$$

$$+ \sum_i \left\{ \sum_{f_j \text{ s.t. } v_i \in f_j} \theta(v_i, f_j) - 2 \sum_{j=1,\ldots,m(i)} \theta(v_i, \tilde{f}_{i,j}) \right\} \quad (21)$$

where $\theta(v, f)$ once again denotes the angle in face f at vertex v. As in our discussion of polygonal curve denoising in Section 3.3, these angles are easily expressed in terms of vertex coordinates using inner products:

$$\theta(v_i, f_j) = \cos^{-1} \left(\frac{(v_{\ell_1(i,j)} - v_i) \cdot (v_{\ell_2(i,j)} - v_i)}{|v_{\ell_1(i,j)} - v_i| \, |v_{\ell_2(i,j)} - v_i|} \right)$$

and there is of course a completely analogous expression for $\theta(v_i, \tilde{f}_{i,j})$. These expressions can now be differentiated with respect to v_i and then used in differentiating (21); the resulting formulae are not unlike those in Section 3.3, such as (13). Using them and an estimate of the area element A_i associated

23

with the i-th vertex v_i obtained from the triangular faces surrounding v_i as in [26], gradient descent can be implemented in a straight forward manner and leads to an ODE system for the $v_i(t)$.

We can look at some simple examples to get an idea about the dependence of (21) on the v_i. Figure (7) shows the dependence of the absolute Gaussian curvature at the vertex of a <u>cone</u> with a circular base on the cone height h. The exact value is $2\pi(1 - \sqrt{\frac{1}{h^2+1}})$, which depends smoothly on h, including at $h = 0$. Figure (8) shows the absolute Gaussian curvature at the center vertex of a surface patch as the coordinates of some of the peripheral vertices are perturbed by a one parameter perturbation. In the process, the surface patch changes "type" in the sense that the center vertex starts out on the surface of the convex hull of its star, but as the perturbation amplitude gets larger, ends up in the interior.

Finally, we give examples of surface fairing with our proposed model. Figure 9 shows denoising results on a man made object (fandisk) whose noise-free version has sharp features such as ridges and corners. The proposed model diffuses out the noise but maintains these features very effectively, as expected. For comparison purposes, the result of denoising the noisy surface via mean curvature motion is also included; it was implemented in the triangulated surfaces setting in essentially the same way as described in [15]. As is well-known, this flow diffuses out the sharp features along with the noise, as can be seen. Figure 10 shows results of the proposed model on another surface that contains some sharp features.

4   Conclusion

The geometry processing model presented in this paper constitutes the natural analogue of the total variation model from image processing; it is based on the *total absolute Gaussian curvature* of a surface. As far as we know, gradient descent for this energy, which leads to an interesting geometric flow of the surface, had not been considered before. It can be used to denoise surfaces while maintaining sharp essential features such as corners and ridges. Drawing on some of the previous work of geometers, we present some first steps towards effective numerical implementation of this flow. Numerical experiments in fairing surfaces with the proposed model yield very promising results.
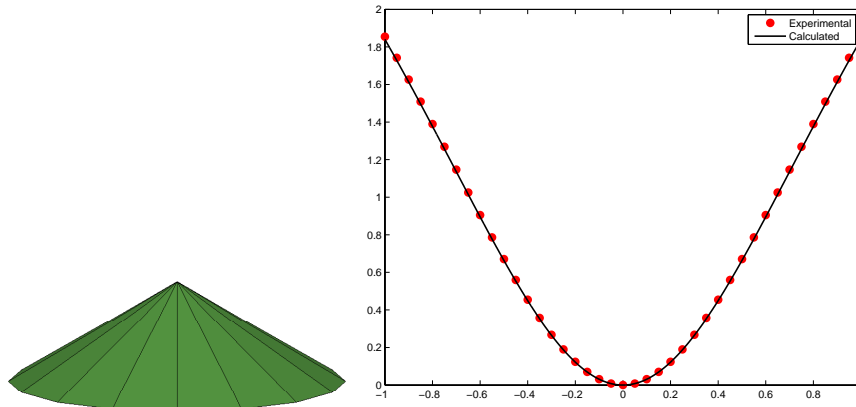
Figure 7: Discrete absolute Gaussian curvature of a surface patch under perturbations to the position of a vertex. Starting from a planar configuration, the central vertex shown on the left figure is pushed up (or down), forming a cone with a circular base. The plot on the right shows the exact and computed values of the absolute value of curvature integrated in a neighborhood of the tip of the cone as a function of its height. The dependence is smooth, as shown.

References

[1] L. Alboul and R. van Damme. Polyhedral metrics in surface reconstruction: Tight triangulations. In T. Goodman and R. Martin, editors, *The Mathematics of Surfaces VIII*, pages 309–336. Clarendon Press, Oxford, 1997.

[2] M. Alexa. Wiener filtering of meshes. In *Proceedings of the Shape Modeling International*, 2002.

[3] T. F. Banchoff. Critical points and curvature for embedded polyhedral surfaces. *Journal of Differential Geometry*, 3:257–268, 1967.

[4] T. F. Banchoff and W. Kühnel. *Tight submanifolds, smooth and polyhedral*, pages 51–118. Cambridge University Press, 1997.

[5] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In Kurt Akeley, editor, *Siggraph 2000, Computer Graphics Proceedings*, pages 417–424. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.
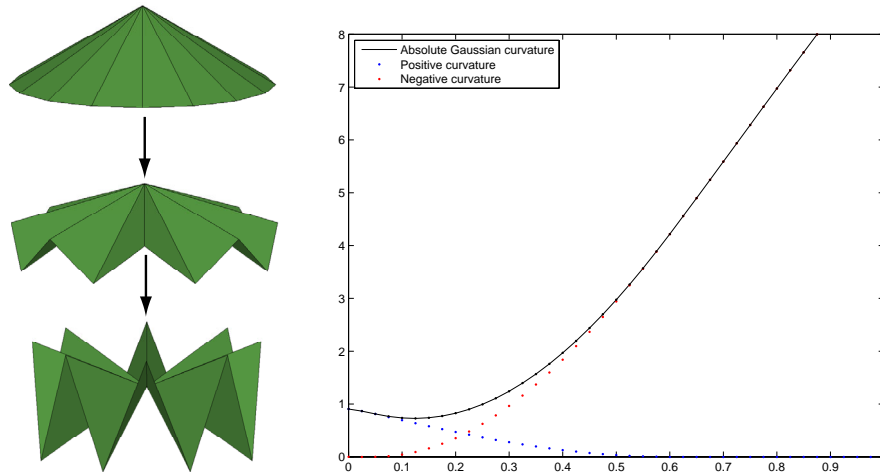
Figure 8: Discrete Gaussian curvature of a surface patch under a one-parameter perturbation to the positions of some of the vertices. As the amplitude of the perturbation grows, the surface patch changes type so that the center vertex is initially on the surface of the convex hull of its star, but eventually ends up in its interior. The plot on the right shows the discrete absolute Gaussian curvature at the center vertex as a function of perturbation amplitude, as well as the positive and negative curvatures.

[6] A. I. Bobenko and P. Schröder. Discrete Willmore flow. In M. Desbrun and H. Pottman, editors, *Eurographics Symposium on Geometry Processing*, pages 101–110, 2005.

[7] M. Buchin and J. Giesen. Minimizing the total absolute Gaussian curvature in a terrain is hard. In *Proceedings of the 17th Canadian conference on computational geometry*, pages 192–195, 2005.

[8] A. Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*, 20:89–97, 2004.

[9] A. Chambolle and P.-L. Lions. Image recovery via total variation minimization and related problems. *Numerische Mathematik*, 76:167–188, 1997.

[10] T. Chan, J. Shen., and S. Kang. Euler's elastica and curvature-based image inpainting. *SIAM Journal on Applied Mathematics*, 63(2):564–592, 2002.

[11] T. F. Chan, G. Golub, and P. Mulet. A nonlinear primal-dual method for total variation based image restoration. *SIAM Journal on Scientific Computing*, 20:1964–1977, 1999.

[12] L.-T. Cheng. Construction of shapes arising from the Minkowski problem using a level set approach. *J. Sci. Comput.*, 19:123–138, 2003.

[13] U. Clarenz, U. Diewald, G. Dziuk, M. Rumpf, and R. Rusu. A finite element method for surface restoration with smooth boundary conditions. *Computer Aided Geometric Design*, 21:427–445, 2004.

[14] U. Clarenz, U. Diewald, and M. Rumpf. Anisotropic diffusion in surface processing. In T. Ertl, B. Hamann, and A. Varshney, editors, *Proceedings of IEEE Visualization 2000*, pages 397–405, 2000.

[15] M. Desbrun, M. Meyer, P. Schröder, and A. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. ACM SIGGRAPH, 1999.

[16] M. Desbrun, M. Meyer, P. Schröder, and A. Barr. Anisotropic feature-preserving denoising of height fields and bivariate data. *Graphics Interface*, 2000.

[17] M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.

[18] M. Droske and M. Rumpf. A level set formulation of willmore flow. *Interfaces and Free Boundaries*, 6:361–378, 2004.

[19] S. Esedoḡlu and J. Shen. Digital inpainting based on the Mumford-Shah-Euler image model. *European Journal of Applied Mathematics*, 13:353–370, 2002.

[20] S. Fleishman, I. Drori, and D. Cohen-Or. Bilateral mesh denoising. In *ACM SIGGRAPH 2003 Proceedings*. ACM SIGGRAPH, 2003.

[21] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Stuetzle. Piecewise smooth surface reconstruction. pages 295–302. ACM SIGGRAPH, 1994.

[22] T. Hou, J. Lowengrub, and M. Shelley. Removing the stiffness from interfacial flows with surface tension. *Journal of Computational Physics*, 114:312, 1994.

[23] Liana Lorigo, Olivier Faugeras, Eric Grimson, Renaud Kerriven, Ron Kikinis, Arya Nabavi, and Carl-Fredrik Westin. Co-dimension 2 geodesic active contours for the segmentation of tubular structures. In *Proceedings of Computer Vision and Pattern Recognition*, 2000.

[24] C. Lu, Y. Cao, and D. Mumford. Surface evolution under curvature flows. *Journal of Visual Communication and Image Representation*, 13:65–81, 2002.

[25] S. Masnou and J.-M. Morel. Level-lines based disocclusion. In *5th IEEE International Conference on Image Processing*, 1998.

[26] M. Meyer, M. Desbrun, P. Schröder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. Proceedings of Visualization and Mathematics, 2002.

[27] H. P. Moreton and C. H. Sequin. Functional optimization for fair surface design. In *Proceedings of SIGGRAPH*, pages 167–176, 1992.

[28] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42:577–685, 1989.

[29] M. Nitzberg, D. Mumford, and T. Shiota. *Filtering, Segmentation, and Depth*. Number 662 in Lecture Notes in Computer Science. Springer-Verlag, 1993.

[30] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulation. *Journal of Computational Physics*, 79:12–49, 1988.

[31] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.

[32] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259 – 268, 1992.

[33] J. Shah. Elastica with hinges. *Journal of Visual Communication and Image Representation*, 13:36–43, 2002.

[34] T. Tasdizen, R. Whitaker, P. Burchard, and S. Osher. Geometric surface processing via normal maps. *ACM Transactions on Graphics*, 22:1012–1033, 2003.

[35] T. Tasdizen, R. T. Whitaker, P. Burchard, and S. Osher. Geometric surface smoothing via anisotropic diffusion of normals. In *Proceedings of IEEE Visualization 2002*, pages 125–132, 2002.

[36] G. Taubin. Curve and surface smoothing without shrinkage. Fifth International Conference on Computer Vision (ICCV'95), 1995. (URL = http://www.research.ibm.com/people/t/taubin).

[37] G. Taubin. A signal processing approach to fair surface design. pages 351–358. SIGGRAPH'95, 1995. (URL = http://www.research.ibm.com/people/t/taubin).

[38] G. Taubin. Linear anisotropic mesh filters. Technical Report RC22213, IBM Research Division, 2001. (URL = http://www.research.ibm.com/people/t/taubin).

[39] W. Welch and A. Witkin. Free-form shape design using triangulated surfaces. In *Proceedings of SIGGRAPH*, pages 247–256, 1994.

[40] Ross. T. Whitaker. Volumetric deformable models: Active blobs. In *Visualization in Biomedical Computing*. SPIE, 1994.

[41] W. Zhu, T. F. Chan, and S. Esedoḡlu. Segmentation with depth: A level set approach. *SIAM Journal on Scientific Computing*, 28:1957–1973, 2006.
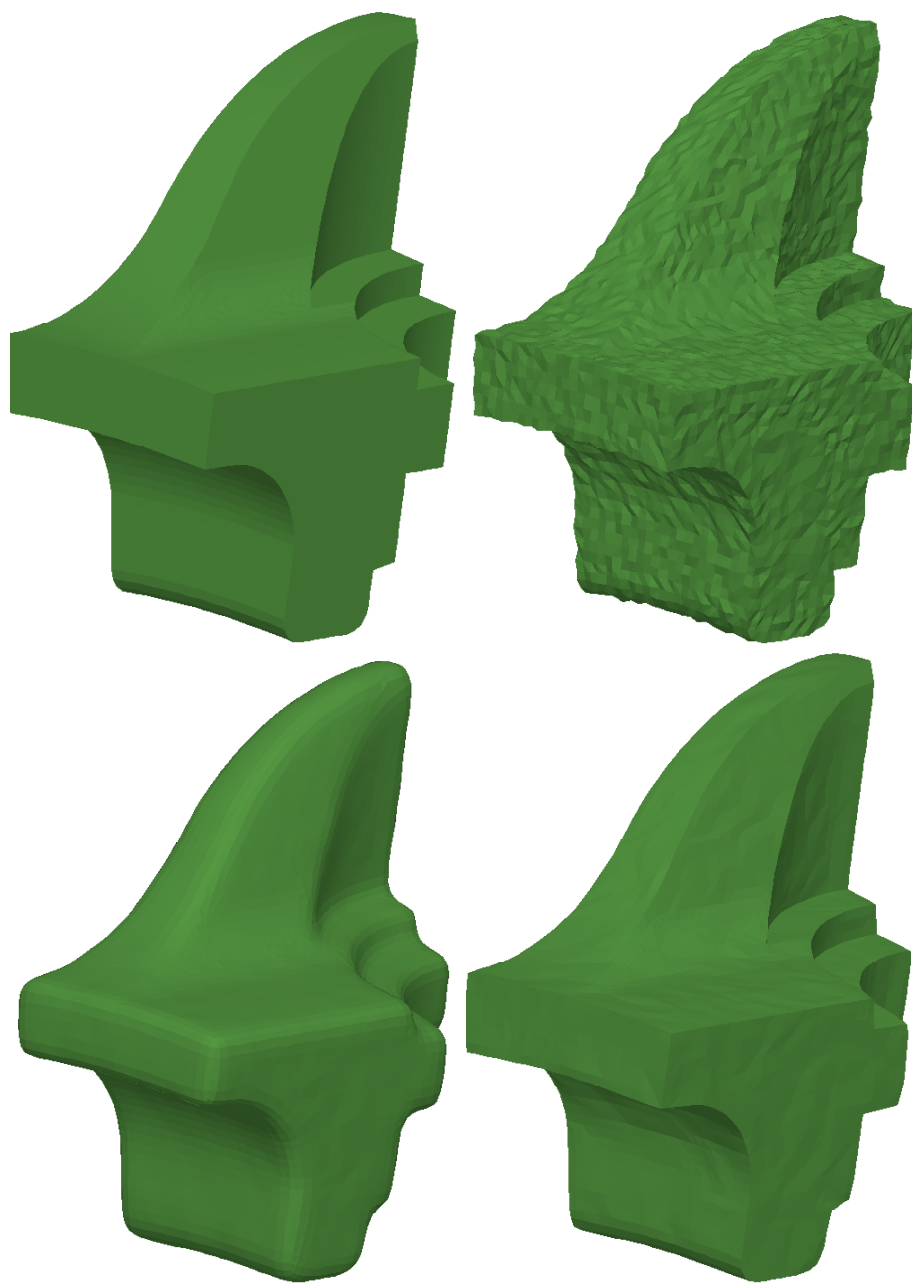
Figure 9: Surface denoising by the proposed model. Upper left surface is the original. Upper right is the its noised up version. Lower left surface is the denoising obtained via motion by mean curvature. Lower right is the result of the proposed algorithm; its evolution is very slow by this stage.
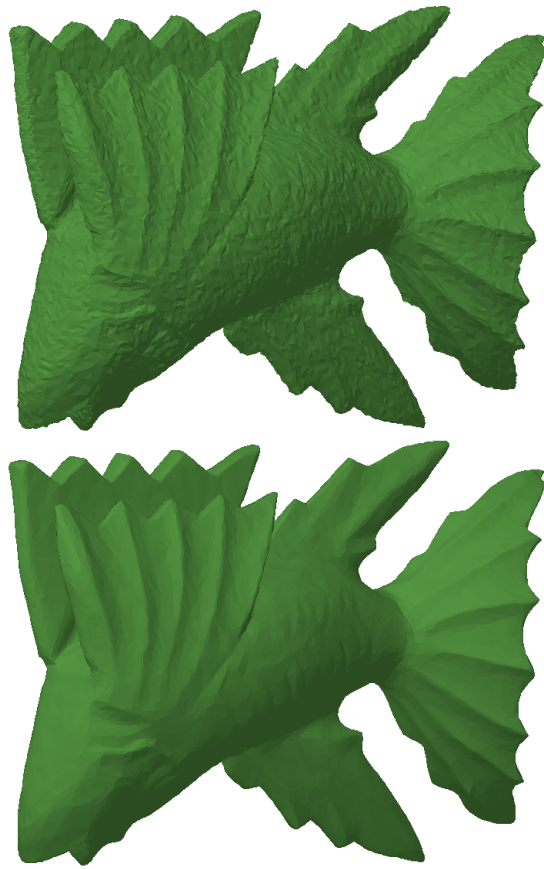
Figure 10: Another example of surface denoising by the proposed model. On the top is the noisy surface. On the bottom is the result of the proposed algorithm.