Pacific Graphics 2015
N. J. Mitra, J. Stam, and K. Xu
(Guest Editors)

Volume 34 (2015), Number 7

Mesh Denoising using Extended ROF Model with L_1 Fidelity

Xiaoqun Wu[1,2] and Jianmin Zheng[2] and Yiyu Cai[2] and Chi-Wing Fu[2]

[1]Beijing Key Lab of Big Data Technology for Food Safety & School of Computer and Information Engineering, Beijing Technology and Business University, Beijing, China
[2]College of Engineering, Nanyang Technological University, Singapore

Abstract

This paper presents a variational algorithm for feature-preserved mesh denoising. At the heart of the algorithm is a novel variational model composed of three components: fidelity, regularization and fairness, which are specifically designed to have their intuitive roles. In particular, the fidelity is formulated as an L_1 data term, which makes the regularization process be less dependent on the exact value of outliers and noise. The regularization is formulated as the total absolute edge-lengthed supplementary angle of the dihedral angle, making the model capable of reconstructing meshes with sharp features. In addition, an augmented Lagrange method is provided to efficiently solve the proposed variational model. Compared to the prior art, the new algorithm has crucial advantages in handling large scale noise, noise along random directions, and different kinds of noise, including random impulsive noise, even in the presence of sharp features. Both visual and quantitative evaluation demonstrates the superiority of the new algorithm.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

1. Introduction

This paper considers the problem of removing noise and outliers from triangular meshes while preserving their underlying surfaces, in particular the fine features. Such a mesh denoising process is necessary and important because mesh models obtained from various acquisition devices such as 3D scanners or reconstructed from data of other formats such as volume data or multi-view images often contain noise and outliers [BAK*10]. Robust and intelligent mesh denoising benefits many digital geometric processing applications.

However, mesh denoising is not trivial in several aspects:

(1) Since both noise and geometric features are of high frequency, it is difficult to distinguish features from noise. Thus, the co-presence of noise/outliers and geometric features such as shape edges and corners requires some intelligent strategies in order to preserve the features in the denoising process.
(2) Many mesh denoising algorithms are directly adapted from their counterparts in image processing. Unlike images, meshes often have irregular connectivity and non-uniform sampling. This implies that special care should be taken in developing the mesh denoising algorithms.
(3) In practice, different types of noise such as Gaussian noise and random impulsive noise may co-exist in a mesh model. Moreover, the noise may corrupt the surfaces along random directions.
(4) Large scale noise could make the situation even worse.

Over the past two decades, extensive research has been done on mesh denoising [Tau95,DMSB99,JDD03,FDCO03, SRML07,ZFAT11,WYL*14,ZWZD15,WYP*15]. Great effort has been devoted to overcoming the first two issues listed above and a wide variety of mesh denoising algorithms have been developed. These algorithms perform well for feature preservation in the presence of small scale noise. However, issues (3) and (4) still challenge the state-of-the-art. For instance, many mesh denoising algorithms implicitly assume that the noise is Gaussian and corrupts the surface along its normal direction. When a mesh has large scale noise or contains random impulsive noise, they usually do not work well.

This paper explores the possibility of overcoming these challenges using a global variational approach. We propose

a new variational model for mesh denoising with three terms: (1) a non-smooth *data-fidelity* term for constraining the solution to approximate the input mesh, (2) a *regularization* term for disfavoring (geometric) oscillations but allowing for the reconstruction of meshes with sharp features, and (3) a *fairness* term for smoothing the solution and avoiding foldovers. It is a natural extension, in the context of digital geometry processing, of the ROF (Rudin, Osher and Fatemi [ROF92]) model with L_1 fidelity [Nik04, Cg05], which is a variational model composed of a fidelity term and a total variation regularization term, and has achieved great success in image denoising and reconstruction. Our work is motivated by two key observations behind the ROF model with L_1 fidelity: (1) reducing the total variation of the signal subject to a close fit to the original one will remove unwanted (geometric) oscillations while preserving important non-smooth features; and (2) the L_1 data-fidelity term encourages exact fitting of the uncorrupted data while replacing the outliers with less dependence on their exact value. As an ROF model with L_1 fidelity, our model has the following properties:

- The L_1 data-fidelity term makes the regularization be less dependent on the exact value of outliers and noise. Thus, our model can properly handle outliers, large scale noise, impulsive noise and uncorrupted data simultaneously.
- The specifically-designed regularization term has no bias against corners and creases. It also avoids the "over-flatten" and "over-sharpen" effects for small scale geometric features, especially in the presence of large scale noise, by L_0 minimization [HS13]. Thus our model is able to preserve geometric features.

These nice properties have been confirmed by extensive experiments.

The contribution of this paper is twofold: (1) we present an extended ROF model with L_1 fidelity for mesh denoising, where the regularization term is formulated as the total absolute edge-lengthed supplementary angle of the dihedral angle. It yields more robust and accurate results than the prior art, especially in the presence of large scale noise, outliers and random impulsive noise; and (2) we provide an augmented Lagrangian method for solving the proposed variational model, which is shown to be efficient and effective.

2. Related Work

There are a large number of mesh denoising algorithms in literature. They can be classified into isotropic approaches, anisotropic diffusion, bilateral filtering, face-normal-based approaches, variational approaches, etc. A good reference is referred to [BAK*10]. Here we just briefly review some of the previous methods that are relevant to our work.

Most early mesh denoising methods are isotropic. A well-known one is Taubin's signal processing method [Tau95], which integrates the diffusion equation and is a linear mesh smoothing method. Since isotropic methods do not preserve geometric features, anisotropic approaches are preferred when the surfaces contain sharp features. Various anisotropic diffusion methods are developed [CDR00, DMSB00, BX03, HP04]. Another way to develop anisotropic methods is to use bilateral filtering, which has some connection to diffusion [Bar02]. For example, Fleishman et al. [FDCO03] and Jones et al. [JDD03] adapt image bilateral filtering for mesh denoising. The basic idea is to anisotropically average the nearby vertices weighted by both spatial difference and vertex difference. Some improvements are presented in [SRML07, ZFAT11]. These methods are simple in both concepts and implementation. They work well when the noise is small in scale and along the normal direction of the surfaces. Recently bilateral filtering is generalized to perform feature-preserving smoothing of signals on images, meshes and other domains within a unified framework, which also forms the basis for constructing mean shift schemes to handle a variety of extremely noisy signals [SCBW14].

Observing that face normal variations can better describe the surface variations than vertex variations, some algorithms are developed to first smooth the face normal field and then update the vertex positions to match the smoothed normal field. Since the vertex position update is simple, the research focuses on smoothing the face normals. For example, Shen et al. [SB04] develop a fuzzy vector median filter for normal face smoothing. Sun et al. [SRML07] improve the method by refining both the smoothing of face normals and the update of vertices. Zheng et al. [ZFAT11] apply bilateral filtering to the facet normal field, producing very good results in general. However, face-normal-based methods are basically two-step approach. It is relatively difficult to control the accuracy of the reconstructed surfaces, especially when the noise is large scale and along random directions.

Recently, several variational approaches based on sparse norm have been proposed to remove noise and outliers from meshes. An analogue of the standard ROF model is presented in the context of geometry processing [EE09], which consists of an L_2 fidelity term and a regularization term formulated as the total absolute Gaussian curvature. While in this model the regularization is vertex-centered, our method formulates the regularization in an edge-centered manner, which is more suitable for preserving sharp edges. Weighted L_1 analysis is used to decouple the noise and sharp features in [WYL*14]. The approach mainly contains two phases. The first phase is to estimate the base mesh by global Laplacian smoothing and the second phase is to recover the sharp features from the residual between the base mesh and the input mesh by weighted L_1 norm. He and Schaefer [HS13] propose an L_0 minimization method for piecewise plane models with large scale noise. This method achieves impressive results for sharp feature preservation and piecewise linear face recovery. However, nonlinear smooth faces and small-scale features of the surfaces are usually over-smoothed. It does not work well when the model contains

large scale impulsive noise. Zhang et al. [ZWZD15] propose a total variation method applied to the face normals to overcome the over-flatten issue caused by an L_0 minimization method. However, their method does not work well when the model contains different kinds of noise and may even cause severe foldovers. Our method is similar to these methods in spirit. Since we adopt the ROF model with L_1 fidelity, our method is more robust and outperforms previous methods in the presence of large scale noise and impulsive noise.

3. Variational Formulation

This section first briefly describes the 1D ROF model with L_1 fidelity, which helps understand the principle of our work before we present our extended model for mesh denoising.

3.1. 1D ROF model with L_1 fidelity

Given an input function $f(x), x \in [a,b]$, the ROF model with L_1 fidelity is to exhibit the reconstructed function as the minimizer of the following energy [Nik04]:

$$\lambda \int_a^b |f(x)-u(x)|dx + \int_a^b |u'(x)|dx. \tag{1}$$

This functional is to be minimized over all $u(x)$. The first term is the L_1 fidelity, encouraging the solution $u(x)$ to approximate $f(x)$. The L_1 fidelity reduces the dependence on the exact value of the variation in the regularization process [Cg05]. The second term is the total variation regularization. Reducing the total variation removes unwanted oscillations while maintaining important discontinuities. In fact, the functions with the least total variation are monotone functions. That is, given the boundary conditions $u(a) = A$ and $u(b) = B$, as long as $u(x)$ is monotone, the total variation is always $|B-A|$, so the regularization sees no difference no matter whether $u(x)$ is smooth or discontinuous. This property allows the model to reconstruct discontinuous signal.

3.2. Extended ROF model with L_1 fidelity for meshes

Assume that the input triangular mesh M^0 is represented by $\{U^0, E^0, T^0\}$, where $U^0 = \{v_i^0\}$ is a set of vertices possibly corrupted by noise and outliers, $E^0 = \{e_i^0\}$ is a set of edges with each edge e_i^0 bounded by two vertices $v_{i_1}^0, v_{i_2}^0$ with $i_1 \neq i_2$, $T^0 = \{\triangle v_i^0 v_j^0 v_k^0\}$ with $i \neq j, j \neq k, i \neq k$ is a set of triangles, and the superscript '0' indicates the initial status. Our goal is to remove noise and outliers whilst preserving the discontinuities in the normals of the underlying clean surface such as sharp edges or corners. This can be accomplished if we can find the following decomposition

$$v_i^0 = v_i + \delta_i \,, \tag{2}$$

where v_i are the denoised vertices and δ_i are noise or outliers contained in v_i^0. The denoised mesh $M = \{U, E, T\}$ is then formed by v_i with the same connectivity as M^0. To obtain a unique solution in the decomposition, we have to add further constraints. Similar to the 1D ROF model, we seek the denoised vertices v_i as the minimizer of the following energy, which we call the *extended ROF model with L_1 fidelity*:

$$\lambda E_{fidelity} + E_{regularizer} + \frac{r}{2} E_{fairness} \,, \tag{3}$$

where $E_{fidelity}$ is the fidelity term, $E_{regularizer}$ is the regularization term, $E_{fairness}$ is the fairness term, and λ and r are the fidelity and fairness parameters, respectively, that balance the effects of the three terms. The detailed formulation of the three terms is given below.

3.2.1. Fidelity $E_{fidelity}$

As shown in [Nik04, Cg05] for signal or image processing, the L_1 formulation of the fidelity enhances the ROF model with many nice properties for feature-preserving denoising. Therefore, we define our fidelity term as

$$E_{fidelity} = \sum_{v_i \in U} \|v_i^0 - v_i\|, \tag{4}$$

which is different from most previous approaches that use L_2 norm: $\sum_{v_i \in U} \|v_i^0 - v_i\|^2$. While our fidelity encourages the denoised vertices v_i to be close to the input vertices, it also avoids the excessive influence of the exact value of outliers and large scale noise. It is especially suitable for removing impulsive noise. Our experiments show that it is also effective for removing Gaussian noise (see Figs. 3,4,5,9,11) though in image processing L_2 norm is typically used for handling Gaussian noise. To the best of our knowledge, this is the first work that uses a non-smooth fidelity term in mesh noising.

3.2.2. Regularizer $E_{regularizer}$

To remove the noise, we need to design an appropriate term to regularize v_i. We may ask what types of meshes are treated as noise free. Note that in the 1D ROF model, the monotone functions are treated as noise free. The regularization is formulated as a total variation, which assigns all monotone functions, whether smooth or discontinuous, the same energy. In the context of mesh denoising, the analogue of monotone functions could be seen as shapes that do not possess unnecessary (geometric) oscillations. This is because a mesh corrupted by noise usually has irregular variation of its normal from face to face, exhibiting (locally frequent) geometric oscillations. Hence, we have to find a quantity to measure the occurrence of oscillations and assign the least possible energy to noise-free shapes.

For the convenience of description, we assume that meshes M^0 and M are polyhedral meshes. Our concept and method can be easily extended to open meshes.

Consider an edge e_i of M. Let θ_i be the interior dihedral angle of M at e_i. If e_i is a convex edge of M, $\theta_i < \pi$. If e_i is a concave edge, $\theta_i > \pi$. We define the supplementary angle of θ_i to be $\pi - \theta_i$. It is positive for a convex edge and

negative for a concave edge, so its sign indicates the relative orientation of two adjacent triangular faces.

Denote by l_i the length of edge e_i. We multiply the supplementary angle by l_i, which gives the total mean curvature of the mesh for the edge [Sul06]. Based on Schläfli formula [Pak10], $\sum_{e_i \in E} l_i(\pi - \theta_i)$ is a constant for a continuous isometric deformation of a polyhedron. Due to the constraint of the fidelity term, we assume that the lengths of the corresponding edges of M and M^0 are roughly the same. Thus, M is approximately an isometric deformation of M^0, and $\sum_{e_i \in E} l_i(\pi - \theta_i) \approx \sum_{e_i^0 \in E^0} l_i^0(\pi - \theta_i^0)$. This suggests that we define our regularization as

$$E_{regularizer} = \sum_{e_i \in E} |l_i(\pi - \theta_i)| \,, \tag{5}$$

which we call the *total absolute edge-lengthed supplementary angle of the dihedral angle* (shortened as "TESA"). It is easy to verify that the TESA has the following properties:

- $\sum_{e_i \in E} |l_i(\pi - \theta_i)| \geq |\sum_{e_i \in E} l_i(\pi - \theta_i)|$
- $\sum_{e_i \in E} |l_i(\pi - \theta_i)| = |\sum_{e_i \in E} l_i(\pi - \theta_i)|$ if and only if all edges are convex or all edges are concave.

Thus, the closer $\sum_{e_i \in E} |l_i(\pi - \theta_i)|$ is to $|\sum_{e_i \in E} l_i(\pi - \theta_i)|$, the closer the distribution of the sign of $l_i(\pi - \theta_i)$ is to all positive or all negative, which implies that minimizing $E_{regularizer}$ has an effect of reducing the occurrence of geometric oscillations. Meanwhile, it can also be seen that $E_{regularizer}$ makes no or less distinction between convex shapes. That is, it does not prefer one to the other, and in particular has no bias against corners and creases. It thus allows the preservation of sharp edges and corners.

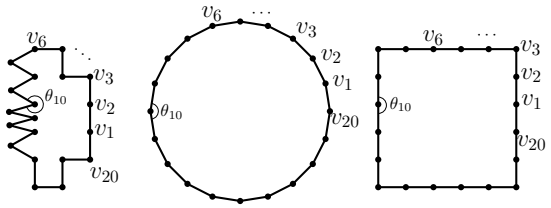

Figure 1: *2D illustration.*

Fig. 1 shows 2D examples to illustrate our regularization, which favors continuity in directional changes. In Fig. 1(left), $\sum_i |(\pi - \theta_i)| \approx 11.32\pi$ is much greater than 2π, which implies that when we go through edges in the polygon (v_1, ..., v_{20}, v_1), we frequently change our heading direction between clockwise and counterclockwise. Our regularizer favors Fig. 1(mid. & right), where $\sum_i |(\pi - \theta_i)| = |\sum_i (\pi - \theta_i)| = 2\pi$ reaches the minimum and the directional change is consistently counterclockwise. Moreover, our regularizer makes no distinction between the circle and square in Fig. 1(mid. & right), enabling the preservation of sharp corners in square.

3.2.3. Fairness $E_{fairness}$

Note that denoising and fairing are actually two different concepts. As pointed out in [FDCO03], mesh fairing alters meshes to increase their degree of fairness while mesh denoising is to remove the noise from meshes. Our fidelity and regularization terms are mainly for the purpose of denoising. Thus, a denoised mesh may not look very smooth. In order to improve the fairness of the denoised mesh, we add a fairness term. Another purpose of adding the fairness term is to reduce the chance of foldover, especially when the noise is large scale and along random directions.

Consider an edge e shared by two triangles $\triangle v_1 v_2 v_3$ and $\triangle v_1 v_3 v_4$. These two triangles can be viewed as an approximation to a bilinear Bézier surface $S(u,v), u, v \in [0,1]$, whose four control points are vertices v_1, v_2, v_3 and v_4. One common practice to fair the surface is to minimize its thin-plate energy:

$$\int_0^1 \int_0^1 \left(S_{uu}^2 + 2S_{uv}^2 + S_{vv}^2 \right) dudv = 2\|v_1 - v_2 + v_3 - v_4\|^2.$$

Thus, we define our fairness term as follows:

$$E_{fairness} = \sum_{e \in E} \|v_1 - v_2 + v_3 - v_4\|^2 \,, \tag{6}$$

which happens to be the same as the term introduced in [HS13].

Remark. Our proposed variational formulation is similar to those proposed in [HS13] and [ZWZD15] in spirit. However, there are two major differences. First, our fidelity term is formulated using an L_1 norm while [HS13] and [ZWZD15] use the conventional L_2 norm. Second, our regularization is formulated using TESA for disfavoring geometric oscillations but allowing for reconstruction of shapes with sharp features. The regularization of [HS13] applies the L_0 norm to an edge-based differential operator to maximize the flat regions of the model, and the regularization of [ZWZD15] is a weighted vectorial total variation of the face normals, which takes the sparsity of sharp features on a clean mesh into consideration. Hence, our formulation may provide relatively simpler and more intuitive geometric explanation of why the sharp features can be preserved in these denoising methods, especially in handling shapes like Fig. 1(middle) and (right).

4. Algorithm

This section describes how to solve the variational model proposed in the preceding section. For this purpose, we need to compute the edge-lengthed supplementary angles. Here, we first give another geometric interpretation of these angles. Referring to Fig. 2, consider an edge $v_1 v_3$ with two opposite vertices v_2 and v_4. Let T_1, T_2 be two vectors of length

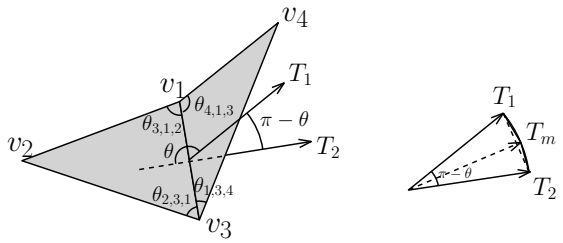

Figure 2: *Notations for the local (3D) configuration of an edge and its supplementary angle.*

$\|v_1v_3\|$, perpendicular to edge v_1v_3. T_1 is the inward facing normal lying on $\triangle v_1v_3v_4$ and T_2 is the outward facing normal lying on the plane of $\triangle v_1v_2v_3$. Then $\pi-\theta$ is the angle between T_1 and T_2, and $\|v_1v_3\|\,|\pi-\theta|$ is the length of the arc that joins the terminal points of T_1 and T_2 while centering at the common initial point of T_1 and T_2 (see Fig. 2(right)).

Rather than directly computing the arc length, which is very complicated, we take an approximation approach. As shown in [HS13], we have

$$\begin{aligned} T_1 &= \cot(\theta_{4,1,3})(v_4-v_3)+\cot(\theta_{1,3,4})(v_4-v_1),\\ T_2 &= \cot(\theta_{2,3,1})(v_1-v_2)+\cot(\theta_{3,1,2})(v_3-v_2). \end{aligned}$$

We compute vector $T_m=\frac{T_1+T_2}{\|T_1+T_2\|}l$. Refer to Fig. 2 (right). T_m is a vector starting at the initial point of T_1 and T_2, and terminating at the middle point of arc $\overset{\frown}{T_1T_2}$. Then we approximate the arc-length by the sum of chord-lengths $|T_1-T_m|+|T_m-T_2|=2|T_m-T_2|$. After some calculations, we obtain $2(T_m-T_2)=$

$$\begin{bmatrix} -\alpha\cot(\theta_{1,3,4})-(2-\alpha)\cot(\theta_{2,3,1}) \\ (2-\alpha)\cot(\theta_{2,3,1})+(2-\alpha)\cot(\theta_{3,1,2}) \\ -\alpha\cot(\theta_{4,1,3})-(2-\alpha)\cot(\theta_{3,1,2}) \\ \alpha\cot(\theta_{4,1,3})+\alpha\cot(\theta_{1,3,4}) \end{bmatrix}^T \begin{bmatrix} v_1\\ v_2\\ v_3\\ v_4 \end{bmatrix} \tag{7}$$

with $\alpha=\frac{2l}{\|T_1+T_2\|}$.

Now we define column vectors $\mathbf{v}=[v_1,v_2,...]^T$ and $\mathbf{v}^0=[v_1^0,v_2^0,...]^T$. Then for all edges of mesh M, their $2(T_m-T_2)$ can be represented as a single expression $\mathbf{K}_1\mathbf{v}$, where $\mathbf{K}_1$ is a matrix whose entries are from the matrix in (7); similarly, all of the terms $v_1-v_2+v_3-v_4$ can be represented as a single expression $\mathbf{K}_2\mathbf{v}$, where $\mathbf{K}_2$ is a constant matrix. Hence, the minimization of (3) can be rewritten in vector form:

$$\min_{\mathbf{v}}\{\lambda\|\mathbf{v}^0-\mathbf{v}\|_1+\|\mathbf{K}_1\mathbf{v}\|_1+\frac{r}{2}\|\mathbf{K}_2\mathbf{v}\|_2^2\}\,, \tag{8}$$

where $\|\cdot\|_1$ and $\|\cdot\|_2$ represent the L_1 and L_2 norms, respectively.

4.1. Augmented Lagrangian method

Since the minimization problem (8) is non-differentiable, we resort the augmented Lagrangian method (ALM) [WZT11] to solve it. For this purpose, two new variables $\mathbf{z}=\mathbf{v}^0-\mathbf{v}$ and $\mathbf{p}=\mathbf{K}_1\mathbf{v}$ are introduced, and the problem is converted to

$$\begin{aligned} &\min_{\mathbf{v},\mathbf{z},\mathbf{p}}\{\lambda\|\mathbf{z}\|_1+\|\mathbf{p}\|_1+\tfrac{r}{2}\|\mathbf{K}_2\mathbf{v}\|_2^2\}\\ &\text{subject to}\quad \mathbf{z}=\mathbf{v}^0-\mathbf{v},\ \ \mathbf{p}=\mathbf{K}_1\mathbf{v}. \end{aligned}$$

We then minimize the following augmented Lagrangian functional

$$\begin{aligned} \mathbf{L}(\mathbf{v},\mathbf{z},\mathbf{p};\lambda_z,\lambda_p)=\lambda\|\mathbf{z}\|_1+\|\mathbf{p}\|_1+\tfrac{r}{2}\|\mathbf{K}_2\mathbf{v}\|_2^2\\ +\langle\lambda_z,\mathbf{z}-(\mathbf{v}^0-\mathbf{v})\rangle+\tfrac{r_z}{2}\|\mathbf{z}-(\mathbf{v}^0-\mathbf{v})\|_2^2\\ +\langle\lambda_p,\mathbf{p}-\mathbf{K}_1\mathbf{v}\rangle+\tfrac{r_p}{2}\|\mathbf{p}-\mathbf{K}_1\mathbf{v}\|_2^2\,, \end{aligned} \tag{9}$$

where λ_z and λ_p are the Lagrangian multipliers, $\langle,\rangle$ is the inner product, and $r_z,r_p>0$ are the penalty coefficients. The entire procedure for solving the problem is outlined in Algorithm 1, which iteratively solves three sub-problems: the **v**-sub problem, **p**-sub problem and **z**-sub problem.

Algorithm 1 ALM algorithm for minimizing (9)

Initialization: $\mathbf{v}^{(0)}=\mathbf{v}^0$, $\mathbf{z}^{(0)}=\mathbf{p}^{(0)}=\lambda_z^{(0)}=\lambda_p^{(0)}=\varepsilon=0$, $k=0$.

while $\varepsilon>\varepsilon_0$ and $k<$ max-iteration-number **do**

Compute $(\mathbf{v}^{(k+1)},\mathbf{z}^{(k+1)},\mathbf{p}^{(k+1)})$ as an (approximate) minimizer of the augmented Lagrangian functional with the Lagrangian multiplier $\lambda_z^{(k)},\lambda_p^{(k)}$:

$$(\mathbf{v}^{(k+1)},\mathbf{z}^{(k+1)},\mathbf{p}^{(k+1)})\approx\arg\min_{\mathbf{v},\mathbf{z},\mathbf{p}}\mathbf{L}(\mathbf{v},\mathbf{z},\mathbf{p};\lambda_z^{(k)},\lambda_p^{(k)})\,,$$

which is done by the following iterations:

Let $\mathbf{z}^{(k+1)}=\mathbf{z}^{(k)}$, $\mathbf{p}^{(k+1)}=\mathbf{p}^{(k)}$.

for $j=1,2,\cdots,K$ **do**

$$\begin{aligned} \mathbf{v}^{(k+1)}&=\arg\min_{\mathbf{v}}L(\mathbf{v},\mathbf{z}^{(k+1)},\mathbf{p}^{(k+1)};\lambda_z^{(k)},\lambda_p^{(k)});\\ \mathbf{z}^{(k+1)}&=\arg\min_{\mathbf{z}}L(\mathbf{v}^{(k+1)},\mathbf{z},\mathbf{p}^{(k+1)};\lambda_z^{(k)},\lambda_p^{(k)});\\ \mathbf{p}^{(k+1)}&=\arg\min_{\mathbf{p}}L(\mathbf{v}^{(k+1)},\mathbf{z}^{(k+1)},\mathbf{p};\lambda_z^{(k)},\lambda_p^{(k)}) \end{aligned}$$

end for

and then update

$$\begin{aligned} \lambda_z^{(k+1)}&=\lambda_z^{(k)}+r_z(\mathbf{z}-(\mathbf{v}^0-\mathbf{v}))\\ \lambda_p^{(k+1)}&=\lambda_p^{(k)}+r_p(\mathbf{p}-\mathbf{K}_1\mathbf{v}) \end{aligned}$$

Update $\mathbf{K}_1$

$\varepsilon=\|\mathbf{v}^{(k+1)}-\mathbf{v}^{(k)}\|_2^2$, k++.

end while

- **v**-sub problem: fix $\mathbf{p},\mathbf{z}$ and solve for $\mathbf{v}$ by minimizing

$$\begin{aligned} \tfrac{r}{2}\|\mathbf{K}_2\mathbf{v}\|_2^2+\langle\lambda_z,\mathbf{z}-(\mathbf{v}^0-\mathbf{v})\rangle+\tfrac{r_z}{2}\|\mathbf{z}-(\mathbf{v}^0-\mathbf{v})\|_2^2\\ +\langle\lambda_p,\mathbf{p}-\mathbf{K}_1\mathbf{v}\rangle+\tfrac{r_p}{2}\|\mathbf{p}-\mathbf{K}_1\mathbf{v}\|_2^2. \end{aligned}$$

This problem can be converted into a sparse linear system.

- **z**-sub problem: fix $\mathbf{v},\mathbf{p}$ and solve for $\mathbf{z}$ by minimizing

$$\lambda\|\mathbf{z}\|_1+\langle\lambda_z,\mathbf{z}-(\mathbf{v}^0-\mathbf{v})\rangle+\frac{r_z}{2}\|\mathbf{z}-(\mathbf{v}^0-\mathbf{v})\|_2^2\,,$$

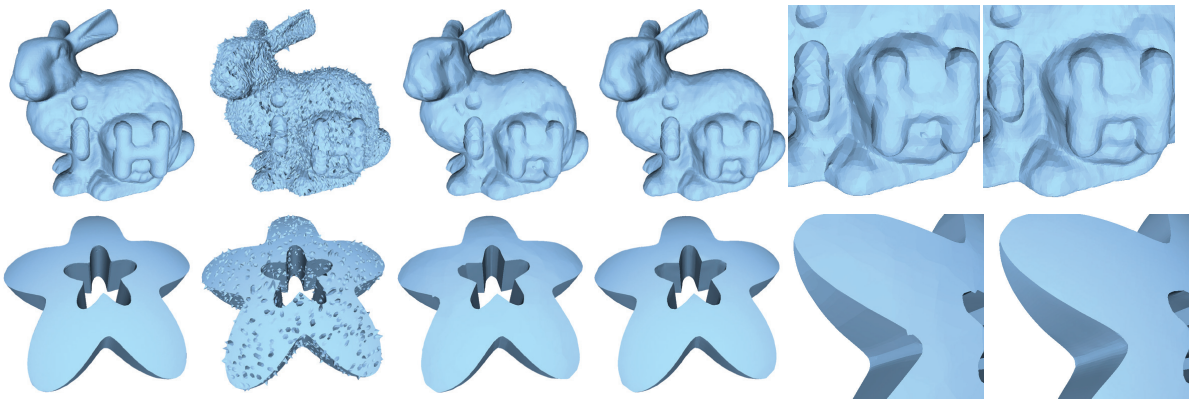

Figure 3: *L_1 fidelity vs L_2 fidelity. From left to right: 1) ground truth, 2) noisy models, 3) denoising results using the L_2 fidelity, 4) denoising results using the L_1 fidelity, 5&6) close-up views of the results in 3) and 4). The model on the top row is corrupted by 10% of impulsive noise with the scale of $[0, l_e]$ and Gaussian noise with $\sigma = 0.15 l_e$ along random directions, while the model on the bottom row is corrupted by 10% of impulsive noise with the scale of $[0, l_e]$, where l_e is the average triangle edge length.*

which has a closed-form solution $\mathbf{z} = \max\{0, 1 - \frac{\lambda}{r_z\|\mathbf{w}_z\|}\}\mathbf{w}_z$ with $\mathbf{w}_z = \mathbf{v}^0 - \mathbf{v} - \frac{\lambda_z}{r_z}$.

- **p**-sub problem: fix $\mathbf{v}, \mathbf{z}$ and solve for $\mathbf{p}$ by minimizing

$$\|\mathbf{p}\|_1 + \langle \lambda_p, \mathbf{p} - \mathbf{K}_1\mathbf{v}\rangle + \frac{r_p}{2}\|\mathbf{p} - \mathbf{K}_1\mathbf{v}\|_2^2 ,$$

which has a closed-form solution $\mathbf{p} = \max\{0, 1 - \frac{1}{r_p\|\mathbf{w}_p\|}\}\mathbf{w}_p$ with $\mathbf{w}_p = \mathbf{K}_1\mathbf{v} - \frac{\lambda_p}{r_p}$.

5. Experimental Results

This section reports our experimental results on a variety of models. Our implementation uses Intel's MKL sparse solver to solve the sparse linear system and the experiments are conducted on Intel(R) Xeon(R) CPU E5-1650. Most of computational cost is taken for solving the linear system. Our current implementation is optimized. The computation time is within 2 minutes for models with fewer than $35k$ vertices.

Like most previous methods, we need to set the parameters properly to produce the best results. Our extended ROF model contains two parameters: λ for controlling the fidelity of the denoised mesh and r for controlling the fairness of the final triangulation. The augmented Lagrangian method also introduces parameters r_z and r_p, which have some denoising effects: the smaller r_z is, the smoother the surface is, whereas the larger r_p is, the more staircase-like the result is. Totally we need to specify four parameter values (λ, r_p, r_z, r) for our method. The parameter sets for other methods are: [FDCO03] (σ_s, vertex iterations), [SRML07] (normal updating iterations, feature detection threshold, vertex updating iterations and neighborhood size), [ZFAT11](local) (normal updating iterations, σ_s, vertex updating iterations, neighborhood size), [ZFAT11](global) (tuning parameter, σ_s, vertex updating iterations, neighborhood size), and [HS13] (smoothness, regularization).

The parameter values we used in our experiments for Figs. 8,9,11 are shown in Table 1. For other examples, the values are given directly in the text. The results generated by [ZWZD15] are kindly provided by the authors.

5.1. L_1 fidelity vs L_2 fidelity

Fig. 3 compares the denoising results by using the L_1 fidelity $\|\mathbf{v} - \mathbf{v}^0\|_1$ and the L_2 fidelity $\|\mathbf{v} - \mathbf{v}^0\|_2^2$. It can be clearly seen that the L_1 fidelity yields better denoising results in removing large scale noise and multiple types of noise. Using L_2, we do not have parameter λ_z. The other three parameter values used here are (0.1, 1000, 100) and (0.01, 1000, 1) for the bunny and star models, respectively. The parameter values for the L_1 fidelity approach are (3,300, 500, 100) and (0.1, 80, 20, 1) for the two models.

5.2. Effects of the fairness term

Fig. 4 shows the effect of the fairness term: it improves the overall surface smoothness and prevents triangle foldover. Without it, some triangles fold over, as indicated in red. Here, the parameter values used are (0.85, 250, 50, 200) and 0 was used to replace 200 for removing the fairness term.

5.3. Visual comparison with other methods

In this subsection, we compare our method with other methods in three different aspects.

1) L_0 minimization. First, we compare our method with L_0 minimization (L_0). In general, both work well for sharp feature preservation, but L_0 over-penalizes the small scale geometry features. Figs. 5 and 6 show the comparison results for two models with small geometry features: the hair area in the Venus head model and the teeth in the Skull model.

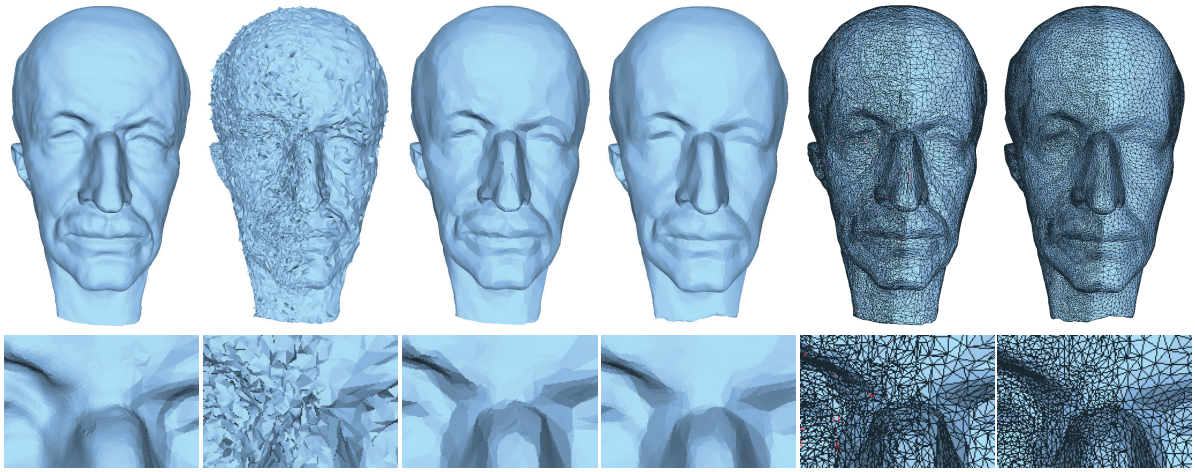

Figure 4: *The effect of the fairness term. From left to right: 1) ground truth, 2) noisy mesh with Gaussian noise ($\sigma = 0.15l_e$) and 10% of impulsive noise with the scale of $[0, l_e]$, 3) denoising result without using the fairness term, 4) denoising result with the fairness term, 5&6) wireframe renderings of the results in 3) and 4). The bottom row shows the close-up views. The red edges in the wireframe renderings (typically shown in the result of the 5th column above) indicate foldovers.*

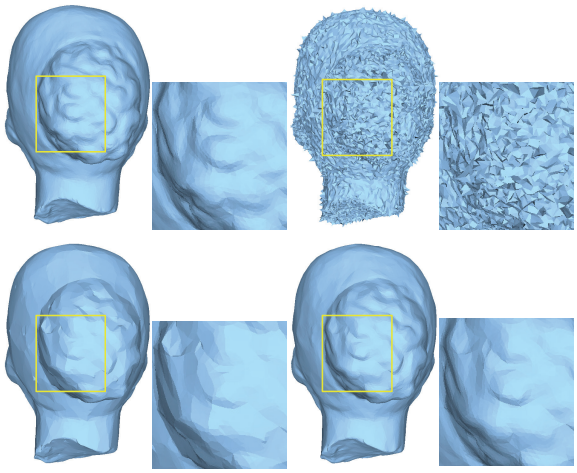

Figure 5: *From left to right, top to bottom: 1) the original Venus head model, 2) the noisy input mesh corrupted by impulsive noise with the scale of $[0, l_e]$ and Gaussian noise with $\sigma = 0.2l_e$ along random directions , 3) the smoothed model by L_0 and 4) the smoothed model by our method.*

In Fig. 5, two types of noise: impulsive noise with the scale of $[0, l_e]$ and Gaussian noise with $\sigma = 0.2l_e$ are added, where l_e is the average triangle edge length. Not only the noise but also the geometry feature like the hair is removed by the L_0 norm. The parameter values are (2, 10) for L_0 and (3, 300, 500, 120) for our method.

Fig. 6 shows the denoising processing of L_0 minimization and our method produced by gradually either increasing the effect of the smoothness term or decreasing the effect of the fidelity term. Specifically, for the L_0 minimization, the regularization parameter value is 1, and the smoothness parameter increases from 1 to 3 and 5 from left to right. For our method, the fidelity parameter λ decreases from 2 to 1 and 0.5 from left to right, and the other three parameters are (500,500,30). To remove the large scale noise near the teeth area, L_0 flattens the small geometry features as shown in the close-up views near the teeth area, while our method not only removes the noise but also preserves the small geometry features well.

Figure 6: *Skull model. From left to right: the second row shows the noisy input and then L_0 norm results by decreasing the fidelity parameter. The third row shows the original mesh and then our results by decreasing the fidelity parameter. The first and fourth rows show the local close-up views.*

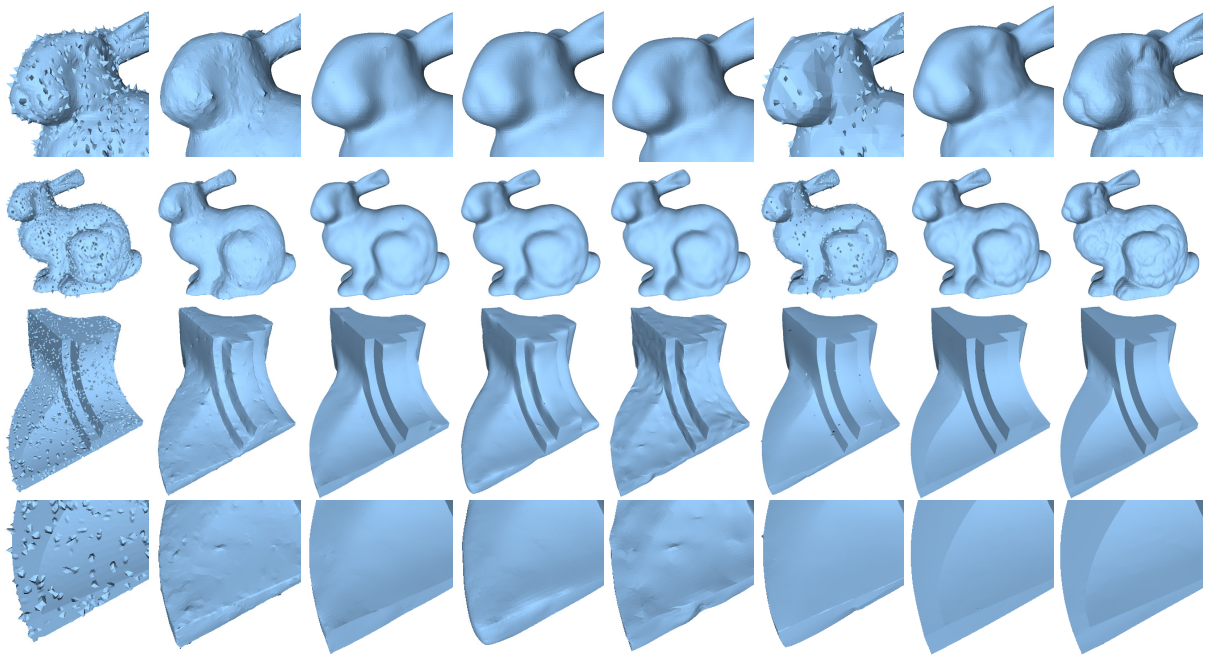

Figure 8: *Comparison with other methods in handling impulsive noise. From left to right: 1) surfaces corrupted by impulsive noise with the scale of* $[0, l_e]$*, denoising results produced by 2) bilateral filtering [FDCO03], 3) face-normal-based denoising [SRML07], 4) bilateral normal denoising (local) [ZFAT11], 5) bilateral normal filtering (global) [ZFAT11], 6)* L_0 *norm [HS13], 7) total variation [ZWZD15], and 8) our method. The top and bottom rows show the local close-up views.*

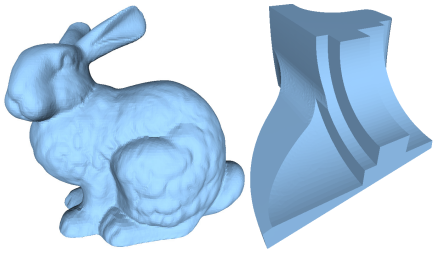

Figure 7: *Ground-truth models: bunny and fandisk.*

2) Impulsive noise. Fig. 8 presents comparison results of handling impulsive noise added to bunny and fandisk, whose ground-truth models are given in Fig. 7. As shown in Fig. 8, our method can effectively remove large scale impulsive noise without over-smoothing the small surface features while the other six methods smooth out the high curvature features in bunny (top) and the shadow edge in Fandisk (bottom). Since most existing methods focus on the design of appropriate smooth filters, they ignore the properties of the fidelity part and thus are not able to filter the large scale noise without removing the model's features. Hence, the over-penalized fidelity on noise will smooth out the surface features, especially for large scale noise. In contrast, our method well preserves the surface features for both the CAD-like and non-CAD models while removing the noise with the new fidelity and regularization terms.

3) Noises along random directions. Fig. 9 (top row) shows that when a model is corrupted by small scale Gaussian noise (with standard deviation $\sigma = 0.1 l_e$) along the normal direction, most previous methods work well, although the first four methods already fail to keep the shadow edge.

If we increase σ to $0.3 l_e$ with noises along random directions, most previous methods fail to remove the noise cleanly, or preserve the surface features well (see the second row of Fig. 9). Meanwhile, severe foldovers are produced. Though [ZWZD15] preserves the surface features well, it cannot avoid the fold-over problem either, as shown in Fig. 10. This may be due to its two-step approach that it first uses variational normal filtering and then updates the vertex position from the modified face normal. When the noise is along random directions, the method may not obtain the correct orientation. Our method works well for small scale noise along normal direction, as well as large scale noise even along random directions.

The bottom row of Fig. 9 presents another comparison, where the fandisk model is simultaneously corrupted by two different types of noise: impulsive noise and Gaussian noise. It can be seen that while other methods have difficulty in removing the noise cleanly, our method can produce high quality results. Fig. 11 performs the same comparison but on a different model.

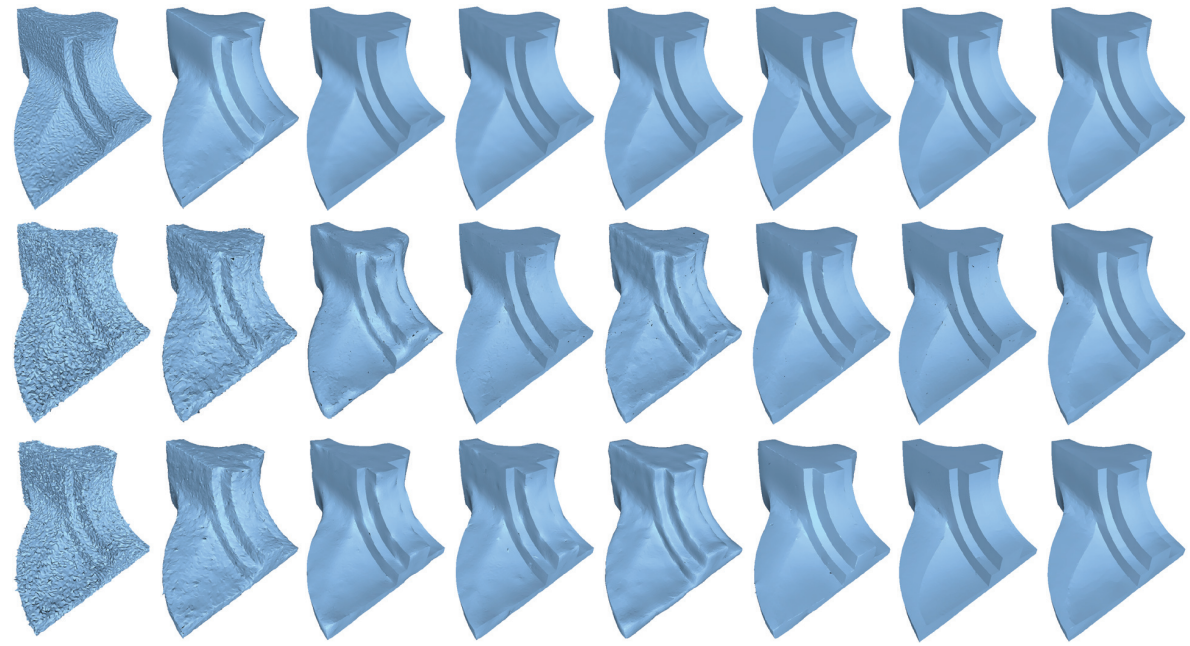

Figure 9: *Comparison with other methods. From left to right: 1) input noisy meshes, results produced by 2) bilateral filtering [FDCO03], 3) face normal denoising [SRML07], 4) bilateral normal denoising (local) [ZFAT11], 5) bilateral normal filtering (global) [ZFAT11], 6) L_0 norm [HS13], 7) total variation [ZWZD15] and 8) our method. From top to bottom: corrupted by Gaussian noise with $\sigma = 0.1l_e$ along normal directions, corrupted by Gaussian noise with $\sigma = 0.3l_e$ along random directions, and corrupted by two types of noise: Gaussian noise with $\sigma = 0.15l_e$ and impulsive noise with the scale of $[0, l_e]$.*

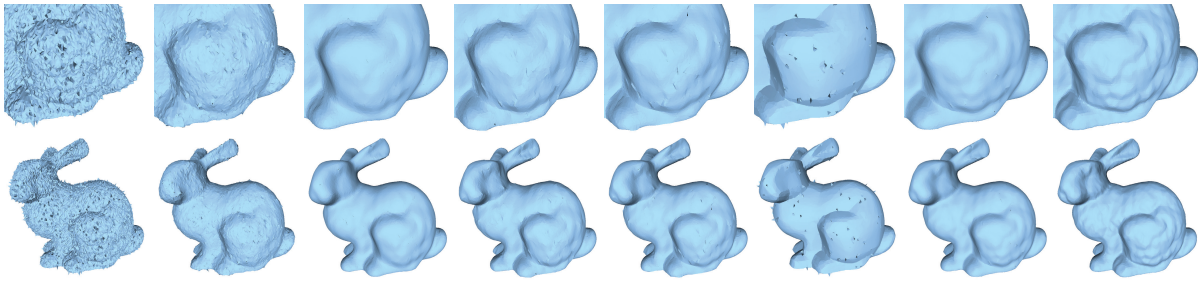

Figure 11: *Comparison with other methods in handling two types of noise simultaneously. From left to right: 1) surface corrupted by Gaussian noise ($\sigma = 0.15l_e$) along random directions and 10% impulsive noise with the scale of $[0, l_e]$, 2) bilateral filtering [FDCO03], 3) face normal denoising [SRML07], 4) bilateral normal denoising (global) [ZFAT11], 5) bilateral normal filtering (local) [ZFAT11], 6) L_0 norm [HS13], 7) total variation [ZWZD15], and 8) our method.*

5.4. Quantitative evaluation

We also conducted quantitative evaluation on our method and previous methods. Three metrics are used to measure the deviation of the denoising results from the ground truth: 1) mean square angular error ($MSAE$) [SRML07], 2) L_2-vertex-based error ($e_{v,2}$), and 3) vertex-based Hausdorff distance (e_H) [ZWZD15]. They measure the angular error, mean error, and maximal error, respectively.

Table 1 lists the results. It can be seen that the denoising results by our method have the least errors in most cases, especially for models containing different types of noise. For Fandisk in Row 1 and Row 2 of Fig. 9, where the noise is small scale Gaussian noise, our method gives the $e_{v,2}$ and e_H values that are not the least, but close to the least.

5.5. Denoising Real Scanned Models

Fig. 12 presents denoising results produced by our method on three different scanned models. In particular, the model on the right is an open mesh. This demonstrates that our method can also work with real data and open meshes.

Table 1: *Quantitative evaluation results: statistics of Fig. 8, 9 and 11, where MSAE, $e_{v,2}$ and e_H are in the scale of $\times 10^{-3}$.*

Methods	Parameters	MSAE	$e_{v,2}$	e_H	Parameters	MSAE	$e_{v,2}$	e_H
	Fig. 8 Row 2, Bunny ($N_v = 34817$, $N_e = 104445$)				Fig. 8 Row 3, Fandisk ($N_v = 25894$, $N_e = 77677$)			
[FDCO03]	(0.3, 10)	129.5	7.3	23.02	(0.5, 10)	105.83	4.7	24.88
[SRML07]	(10, 0.25, 20, NII)	41.62	3.17	25.32	(15, 0.25, 20, NII)	12.69	1.24	9.51
[ZFAT11] (Local)	(10, 0.6, 20, NII)	32.45	2.69	13.12	(20, 0.45, 20, NII)	34.91	2.22	16.6
[ZFAT11] (Global)	(0.05, 0.9, NII)	32.43	2.5	12.05	(0.07, 0.55, 20 NII)	13.04	1.62	13.1
[HS13]	(10, 1)	124.5	5.01	39.47	(200, 1)	38.81	3.02	48.67
[ZWZD15]	N.A.	36.97	2.98	13.12	N.A.	22.08	1.05	5.16
Ours	(2, 500, 500, 100)	**6.64**	**0.6**	**11.03**	(0.1, 180, 20, 1)	**5.4**	**0.77**	**4.93**
	Fig. 9 Row 1, Fandisk ($N_v = 25894$, $N_e = 77677$)				Fig. 9 Row 3, Fandisk ($N_v = 25894$, $N_e = 77677$)			
[FDCO03]	(0.5, 10)	111.5	5.34	19.73	(0.9, 10)	167.24	4.08	24.05
[SRML07]	(5, 0.5, 20, NII)	1.01	0.52	5.42	(15, 0.25, 20, NII)	37.81	1.4	30.17
[ZFAT11] (Local)	(5, 0.3, 20, NII)	0.93	0.47	4.58	(15, 0.4, 20, NII)	40.1	1.42	24.1
[ZFAT11] (Global)	(0.07, 0.3, 20, NII)	0.82	0.4	**3.14**	(0.01, 0.5, 20, NII)	60.86	2.66	14.02
[HS13]	(3, 1)	0.52	0.36	4.48	(100,1)	55.65	3.5	37.17
[ZWZD15]	N.A.	0.32	**0.35**	4.53	N.A.	50.3	3.01	15.89
Ours	(0.1, 150, 200, 1)	**0.31**	0.39	4.01	(0.1, 200, 200, 1)	**9.39**	**1.21**	**13.97**
	Fig. 9 Row 2, Fandisk ($N_v = 25894$, $N_e = 77677$)				Fig. 11, Bunny ($N_v = 34817$, $N_e = 104445$)			
[FDCO03]	(0.9, 10)	367.9	3.58	15.3	(0.6, 10)	129.56	7.3	23.03
[SRML07]	(30, 0.15,20, NII)	468.8	3.63	21.51	(10, 0.2, 20, NII)	39.48	3.18	16.53
[ZFAT11] (Local)	(35, 0.3, 20, NII)	492.6	1.69	11.04	(15,0.4, 20, NII)	44	2.46	28.61
[ZFAT11] (Global)	(0.05, 0.6, 20, NII)	525.5	**1.56**	**8.71**	(0.04, 0.5, 20, NII)	31.96	2.5	13.24
[HS13]	(5, 1)	37.48	2.86	11.19	(30, 1)	80.93	4.38	41.44
[ZWZD15]	N.A.	538.5	3.56	9.87	N.A.	50.3	3.15	15.53
Ours	(0.1, 150, 130, 1)	**19.9**	1.63	9.93	(2, 500, 500, 100)	**8.82**	**0.87**	**12.99**

6. Conclusion

We presented an extended ROF model with L_1 fidelity for feature-preserved mesh denoising. The L_1 data-fidelity term and the regularization term formulated as the total absolute edge-lengthed supplementary angle of dihedral angle make our model be able to remove outliers and different kind of noise (even large scale), while preserving small geometric details and sharp features. The fairness term in our model has the smoothing effect and reduces the occurrence of foldovers. Extensive experiments also showed that our proposed method always outperforms the state-of-the-art.

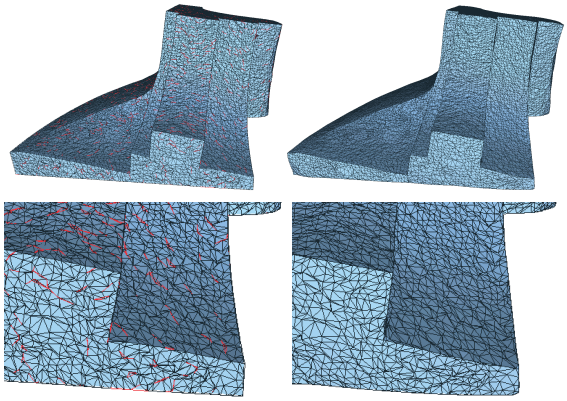

Figure 10: *Denoising results (in wireframe rendering) produced by [ZWZD15] (left) and ours (right). See also Fig. 9 (2nd row). The folded triangles are indicated by red edges.*

Figure 12: *Denoising real scanned models: input models with noise (top) and denoised models (bottom).*

Acknowledgments

We thank anonymous reviewers for the detailed review and constructive comments. We are very grateful to Dr Chunlin Wu for his generous help in our experiments. This research is partially supported by MOE Tier-2 grant (MOE2011-T2-2-041) of Singapore, open funding project of State Key Lab of Virtual Reality Technology and System at Beihang University under Grant NO. BUAA-VR-14KF-04, and Singapore Millennium Foundation which is funded by Temasek Trust.

References

[BAK*10] BOTSCH M., ALLIEZ P., KOBBELT L., PAULY M., LÉVY B.: *Polygon Mesh Processing*. A K Peters, 2010. 1, 2

[Bar02] BARASH D.: A fundamental relationship between bilateral filtering, adaptive smoothing, and the nonlinear diffusion equation. *IEEE Transactions on Pattern Analysis and Machine Intelligence 24* (2002), 844–847. 2

[BX03] BAJAJ C. L., XU G.: Anisotropic diffusion of surfaces and functions on surfaces. *ACM Trans. Graph. 22*, 1 (2003), 4–32. 2

[CDR00] CLARENZ U., DIEWALD U., RUMPF M.: Anisotropic geometric diffusion in surface processing. In *Proceedings of IEEE Visualization 2000* (Washington, DC, USA, 2000), IEEE Computer Society. 2

[Cg05] CHAN T. F., ŕGLU S. E.: Aspects of total variation regularized l_1 function approximation. *SIAM J. Appl. Math* (2005). 2, 3

[DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A. H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *SIGGRAPH '99* (1999), pp. 317–324. 1

[DMSB00] DESBRUN M., MEYER M., SCHRÖDER P., BARR A. H.: Anisotropic feature-preserving denoising of height fields and bivariate data. In *Graphics Interface* (2000), pp. 145–152. 2

[EE09] ELSEY M., ESEDOGLU S.: Analogue of the total variation denoising model in the context of geometry processing. *Multiscale Modeling & Simulation 7*, 4 (2009), 1549–1573. 2

[FDCO03] FLEISHMAN S., DRORI I., COHEN-OR D.: Bilateral mesh denoising. *ACM Trans. Graph. 22*, 3 (July 2003), 950–953. 1, 2, 4, 6, 8, 9, 10

[HP04] HILDEBRANDT K., POLTHIER K.: Anisotropic filtering of non-linear surface features. *Comp. Graph. Forum 23*, 3 (2004), 391–400. 2

[HS13] HE L., SCHAEFER S.: Mesh denoising via L_0 minimization. *ACM Trans. Graph. 32*, 4 (July 2013), 64:1–64:8. 2, 4, 5, 6, 8, 9, 10

[JDD03] JONES T. R., DURAND F., DESBRUN M.: Non-iterative, feature-preserving mesh smoothing. *ACM Trans. Graph. 22*, 3 (July 2003), 943–949. 1, 2

[Nik04] NIKOLOVA M.: A variational approach to remove outliers and impulse noise. *J. Math. Imaging Vis. 20*, 1-2 (2004), 99–120. 2, 3

[Pak10] PAK I.: Lectures on Discrete and Polyhedral Geometry, 2010. 4

[ROF92] RUDIN L. I., OSHER S., FATEMI E.: Nonlinear total variation based noise removal algorithms. *Phys. D 60* (1992), 259–268. 2

[SB04] SHEN Y., BARNER K.: Fuzzy vector median-based surface smoothing. *IEEE Transactions on Visualization & Computer Graphics 10*, 3 (2004), 252–265. 2

[SCBW14] SOLOMON J., CRANE K., BUTSCHER A., WOJTAN C.: A general framework for bilateral and mean shift filtering. *ArXiv e-prints 32* (April 2014). 2

[SRML07] SUN X., ROSIN P. L., MARTIN R. R., LANGBEIN F. C.: Fast and effective feature-preserving mesh denoising. *IEEE Transactions on Visualization & Computer Graphics 13*, 5 (Sept. 2007), 925–938. 1, 2, 6, 8, 9, 10

[Sul06] SULLIVAN J. M.: Curvature measures for discrete surfaces. In *ACM SIGGRAPH 2006 Courses* (New York, NY, USA, 2006), ACM, pp. 10–13. 4

[Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *SIGGRAPH '95* (1995), pp. 351–358. 1, 2

[WYL*14] WANG R., YANG Z., LIU L., DENG J., CHEN F.: Decoupling noise and features via weighted l_1-analysis compressed sensing. *ACM Trans. Graph. 33*, 2 (Apr. 2014), 18:1–18:12. 1, 2

[WYP*15] WEI M., YU J., PANG W.-M., WANG J., QIN J., LIU L., HENG P.-A.: Bi-normal filtering for mesh denoising. *IEEE Transactions on Visualization & Computer Graphics 21*, 1 (Jan 2015), 43–55. 1

[WZT11] WU C., ZHANG J., TAI X.-C.: Augmented Lagrangian method for total variation restoration with non-quadratic fidelity. *Inverse Problems and Imaging 5* (2011), 237–261. 5

[ZFAT11] ZHENG Y., FU H., AU O. K.-C., TAI C.-L.: Bilateral normal filtering for mesh denoising. *IEEE Transactions on Visualization & Computer Graphics 17*, 10 (Oct. 2011), 1521–1530. 1, 2, 6, 8, 9, 10

[ZWZD15] ZHANG H., WU C., ZHANG J., DENG J.: Variational mesh denoising using total variation and piecewise constant function space. *IEEE Transactions on Visualization & Computer Graphics 21*, 7 (2015). 1, 3, 4, 6, 8, 9, 10