Surface Reconstruction from Normals: A Robust DGP-based Discontinuity Preservation Approach

Wuyuan Xie [1], Miaohui Wang [2,⋆], Mingqiang Wei[3], Jianmin Jiang[1], Jing Qin[4]
[1]College of Computer and Software Engineering, Shenzhen University (SZU), P. R. China
[2]Guangdong Key Laboratory of Intelligent Information Processing, College of Information Engineering, SZU, P. R. China
[3]School of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, P. R. China
[4]School of Nursing, Hong Kong Polytechnic University, Hong Kong, P. R. China
{wyxie, mhwang, jianmin.jiang}@szu.edu.cn, mqwei@nuaa.edu.cn, harry.qin@polyu.edu.hk

Abstract

In 3D surface reconstruction from normals, discontinuity preservation is an important but challenging task. However, existing studies fail to address the discontinuous normal maps by enforcing the surface integrability in the continuous domain. This paper introduces a robust approach to preserve the surface discontinuity in the discrete geometry way. Firstly, we design two representative normal incompatibility features and propose an efficient discontinuity detection scheme to determine the splitting pattern for a discrete mesh. Secondly, we model the discontinuity preservation problem as a light-weight energy optimization framework by jointly considering the discontinuity detection and the overall reconstruction error. Lastly, we further shrink the feasible solution space to reduce the complexity based on the prior knowledge. Experiments show that the proposed method achieves the best performance on an extensive 3D dataset compared with the state-of-the-arts in terms of mean angular error and computational complexity.

1. Introduction

Surface reconstruction from normals is driven by several computer vision tasks, such as shape from shading (SfS), deflectometry, and photometric stereo (PS). Given a static scene captured by a camera with fixed viewpoint, PS [15, 16, 20, 21] is widely used to obtain the surface normal orientations. To reconstruct a surface from the estimated normals, the most common way [1, 3, 8, 11, 17] is to conduct integration on the depth gradients which are transformed from the normal orientations. However, the normal map produced by PS may not be uniformly integrable, due to occlusion boundary, sharp surface changes, or computation error. To reconstruct a surface from such a discontinuous and noisy normal map, Terzopoulos *et al.* [18] provided a compact framework to compute the depth field and surface discontinuity simultaneously, where a rough depth map was needed for initialization. Approaches in [6, 7, 9] used binary weights to indicate the presence of surface discontinuities, so that the integrability constraint can be applied partially. Although these methods can reconstruct surfaces from discontinuous and noisy normals, none of them satisfactorily preserves the discontinuity, where the reconstructed surfaces are somewhat distorted. One key reason is that forming a discontinuity between two adjacent patches requires breaking their connection. The aforementioned approaches assume that the number of total reconstructed vertexes is equal to the number of normal orientations, which means every patch is connected, and distortion in the local shape is propagated to the whole surface. Fig. 1 shows an example of reconstructing the boundary vertexes of the *Saddle* model, the object surface is split and can be separately to form a more accurate shape.

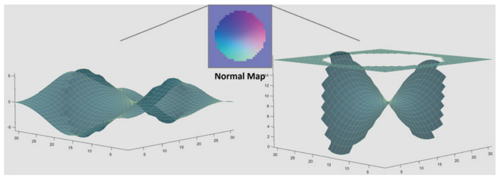

Figure 1. Reconstruction surfaces from a saddle normal map padded with the orientation vectors of $[0, 0, 1]^T$. (Left): Xie *et al.* [24]; (Right): Proposed.

In this paper, we concentrate on addressing the preservation of the discontinuity feature in surface reconstruction

⋆ Corresponding author: Miaohui Wang (wang.miaohui@gmail.com).
The work described in this paper was supported in part by NSFC 61701310, in part by Hong Kong Research Grants Council (Project no. PolyU 152035/17E), and in part by Inlife-Handnet Open Fund 2018 and Start-up Fund of Shenzhen University 2018080. We would like to give thanks to NVIDIA for their generous GPU support.

from normals. The input is a single normal map which can be obtained by PS or point cloud, and the output is a *3D* surface with rich discontinuities. We study and model the normal discontinuity in a discrete geometry domain [24] which is used for surface reconstruction. To measure the incompatibility between two adjacent normal orientations, we introduce two new features, including *absolute depth difference* and *normal angle difference*. The surface discontinuities are detected by these two normal incompatibility features, and followed by vertex partition and redistribution. We build a geometry-based energy function to optimize the discontinuity and the surface shape simultaneously. Besides, the prior statistical knowledge of the normal incompatibility features greatly facilitates obtaining the optimal result with low computation complexity. Extensive experimental results (*e.g.*, comparisons on 16 different models) validate the efficiency of the proposed method.

Comparing to [24], the proposed method accurately reconstructs the surface discontinuities by allowing the mesh node to have multiple vertexes, rather than approximating it in the continuous domain. To the best of our knowledge, this is the first approach to address the discontinuity problem in a discrete geometry way, which has the following main merits:

- **Discontinuity preservation:** As two adjacent facets oriented by incompatible normals are taken to be independent, the discontinuity between them can be formed along the breaking boundary. Meanwhile, as a light-weight least-square optimization is used to compute the depth map, and the surface smoothness on continuous segments is preserved.
- **Robustness to noise:** This research can be applied to the continuous/discontinuous normal maps with noise. Examples with heavy Gaussian noise (*e.g.*, variance up to $8°$ angle degree) can be successfully reconstructed, where discontinuities are still accurately distinguished from noise and preserved.
- **Parameter-tuning free:** Due to the convexity of the proposed surface energy model, the optimal parameter used for discontinuity detection can be uniquely determined. As a result, the proposed approach can be adaptive to varying continuous/discontinuous normals without parameter tuning.

2. Related Work

This paper aims to reconstruct a surface from normals and allow the normal map to contain occlusion boundary, discontinuities, and noise. Traditional approaches [6, 9, 13] enforce the integrability constraint over the whole normal map to estimate the depths, and can be mainly categorized into regularization-based, weighting-based, and basis function-based methods.

Regularization-based: Regularization method attempts to smooth the depth gradients by introducing additional constraints. The most common way [4] is to use the least square energy function which is defined on residuals between the depth difference and the depth gradient. Although it is useful to noisy gradients, its quadratic behavior may over-penalize discontinuities which have higher residuals. Methods in [2, 18] considered the discontinuity through a variational formulation where a rough depth map was needed to estimate discontinuities and depth values simultaneously. Petrovic *et al.* in [11] used a belief propagation to enforce integrability constraints in Markov Random Field network.

Weighting-based: Instead of enforcing the regularization uniformly over the whole gradient field, methods in [6, 7] employed the constraint partially by a weighting map, where the weight values were designed to be inversely proportional to the probability of lying on a discontinuity. An extreme case is Mumford *et al.* [9], where a binary template was used to indicate the presence of discontinuities. Studies in [12, 14, 19, 23] used an expectation maximization method to estimate a weighted discontinuity map, and then reconstructed the surface in segments according to the discontinuity labels. It is worth noting that all these methods either require manually labeling or thresholding for edge detection, thus cannot guarantee the reconstruction quality for a general surface.

Basis function-based: Frankot and Chellappa [3] projected the densely non-integrable gradients onto a set of Fourier basis functions for enforcing surface integrability. Hsieh *et al.* [5] and Karacali *et al.* [7] used wavelet basis functions instead. Kovesi *et al.* [8] computed the correlations of gradients via a bank of shapelets, which could solve the gradients with a certain ambiguity by carefully parameter tuning. One inevitable drawback of all these basis function-based approaches is that the global integrability enforcement constraint on the whole gradient field can smooth out the occlusion boundary. To address this problem, Wu *et al.* in [10, 22] further developed to introduce the kernel basis function, but these approaches are limited to the normal map of rectangular boundary and bring extremely computational complexity.

3. Discontinuity Detection

This research addresses the discontinuity problem based on a *discrete geometry processing* (DGP) framework. In this section, we first give a brief review of the DGP deformation method, and discuss the difficulty of reconstructing the surface discontinuities by the state-of-the-art method in a discrete geometry way. Considering the characteristics of DGP, we introduce two new features to measure the incompatibility between normals, and demonstrate how to use these two features to detect discontinuities.

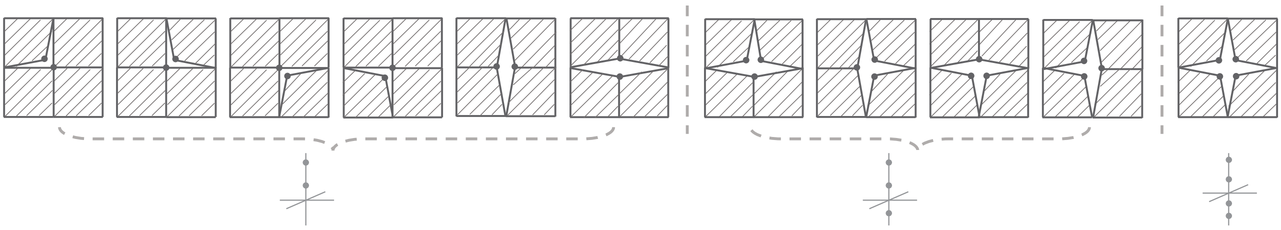

Figure 2. Example of all possible splitting patterns for a vertex with four facets.

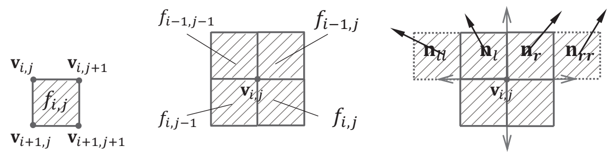

Figure 3. Example of a discrete mesh consisting of vertexes and facets. (Left): A facet and its four vertexes; (Middle): A vertex and the facets around; (Right): Related normals used to calculate the depth difference along the positive direction of the $y-$axis at $\mathbf{v}_{i,j}$.

3.1. DGP-based Shape Deformation

In DGP deformation [24], the surface to be reconstructed is taken as a mesh consisting of a set of micro square facets oriented by the input normal map, and then a local/global deformation is iteratively conducted on the mesh to generate the final shape. Specifically, each pixel (i, j) in a normal map is converted into a facet $f_{i,j}$ bounded by four vertexes $\mathbf{v}_{i,j}$, $\mathbf{v}_{i,j+1}$, $\mathbf{v}_{i+1,j+1}$, and $\mathbf{v}_{i+1,j}$ as shown in Fig. 3 (Left). A vertex $\mathbf{v}_{i,j}$ has its $x-$ and $y-$coordinates fixed, and $z-$coordinate (*i.e.*, the depth value $z_{i,j}$) as an unknown variable to be determined. In each iteration of the shape deformation, a local shaping step is first performed to determine the position of each facet according to its demanding normal orientation. Meanwhile, a global blending step is applied to generate all facets into a connected mesh surface. The local/global steps can be summarized as a least square formulation:

$$\Phi\left(z_{i,j}\right) = \sum_{f_{i,j}} \left\|\mathbf{z}\left(f_{i,j}\right) - \mathbf{p}\left(f_{i,j}\right)\right\|^2, \tag{1}$$

where $\mathbf{z}(f_{i,j})$ is a column vector stacking the depth values of four vertexes of $f_{i,j}$, $\mathbf{z}(f_{i,j}) = [z_{i,j} z_{i,j+1} z_{i+1,j+1} z_{i+1,j}]^T$, and $\mathbf{p}(f_{i,j})$ is the updated depth values of $\mathbf{z}(f_{i,j})$ after employing the local shaping. By iteratively solving Eq. (1), the depth converges to form a surface whose orientations are expected to be the same as the input normals. However, in [24], all facets are supposed to be connected, and the reconstructed shape for the discontinuous normals can be inevitably distorted as shown in Fig. 1.

Actually, for a vertex $\mathbf{v}_{i,j}$ with discontinuity, the

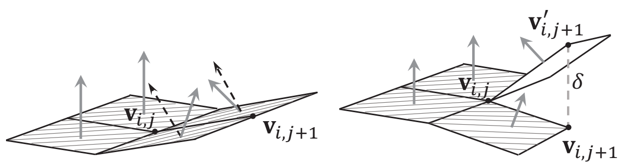

Figure 4. Illustration of the conflicting normals. Arrows in gray are the ground-truth normals, while arrows in black are the reconstructed ones.

$z-$coordinate value is not unique. Let us denote it as $[z_{i,j}^{(1)}, ..., z_{i,j}^{(K)}]^T$. The total number K should be greater than one, and less or equal to the total number of facets around $\mathbf{v}_{i,j}$. We enumerate all possible splittings for $\mathbf{v}_{i,j}$ in Fig. 2, which in general can be grouped into four types: I) $\mathbf{v}_{i,j}$ has a unique depth value, and four facets are connected and continuous; II) $\mathbf{v}_{i,j}$ has two depth values, *e.g.,* the first group separated by dotted line in Fig. 2. There are two cases in this group: 1) Three facets are connected to one split vertex, and the last one independently holds another vertex alone (*i.e.*, the first four graphs); 2) Two pairs of adjacent facets[1] are connected by one split vertex (*i.e.*, the last two graphs of the first group); III) $\mathbf{v}_{i,j}$ has three depth values, *e.g.,* the second group in Fig. 2. A pair of facets is connected to one vertex, and the remaining two facets hold one vertex independently; IV) $\mathbf{v}_{i,j}$ has four depth values, *e.g.,* the third group in Fig. 2, where every facet holds one independent vertex. How to determine the splitting pattern for a vertex $\mathbf{v}_{i,j}$ is the main issue that this research is going to address in the rest of this section.

3.2. Absolute Depth Difference Feature

An intuitive way to detect the compatibility between two adjacent normals is to check whether the facets oriented by them can be connected. Let us consider a patch $\mathcal{P}_{i,j}$ consisting of 2×2 facets and centering at $\mathbf{v}_{i,j}$ as shown in Fig. 3 (Middle). If there are two adjacent conflicting orientations, the associated two facets are divided along the $z-$axis but still hold the joint vertex $\mathbf{v}_{i,j}$ as illustrated in Fig. 4. It is observed that the depth difference between these two split

[1] In this paper, “adjacent” refers to “adjacent but non-diagonal”. We do not consider the discontinuity on the diagonal joints due to its weak connection relationship.

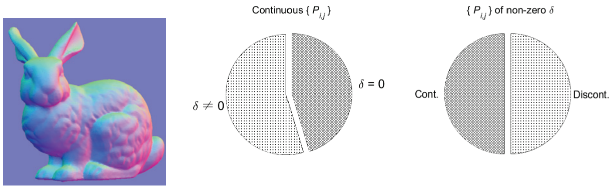

Figure 5. Statistics of the ADD feature for the normal map. (Left): The normal map of *Bunny* padded with $[0, 0, 1]^T$; (Middle): The ADD statistics of all continuous $\mathcal{P}_{i,j}$; (Right): The discontinuity statistics of all $\mathcal{P}_{i,j}$ with the non-zero ADDs.

facets can be used to measure the degree of incompatibility, and defined as the *absolute depth difference* (ADD) feature,

$$\delta_{i,j}^{\ell} = \left| \frac{\boldsymbol{\ell} \cdot \mathbf{n}_l}{\mathbf{k} \cdot \mathbf{n}_r} - \frac{\boldsymbol{\ell} \cdot \mathbf{n}_r}{\mathbf{k} \cdot \mathbf{n}_l} \right|, \tag{2}$$

where $\boldsymbol{\ell}$ has the four directions in $\mathcal{P}_{i,j}$ as shown in Fig. 3 (Right), which is a unit vector of directions of the $x-$ and $y-$ axis in the Cartesian coordinate. $\mathbf{k}$ is a unit vector of the positive direction of the $z-$ axis. $\mathbf{n}_l$ and $\mathbf{n}_r$ are the left and right normal orientations along the direction $\boldsymbol{\ell}$ at $\mathbf{v}_{i,j}$.

Fig. 5 shows the statistics of the ADD feature for all continuous $\mathcal{P}_{i,j}$ from the *Bunny* model. One can see that there is about 60% continuous $\mathcal{P}_{i,j}$ having a non-zero δ. In addition, the distribution of discontinuities for the non-zero δ is also shown in Fig. 5 (Right). For a noise-free discontinuous normal map, all facets are supposed to be connected except for the occlusion boundary. In such a case, the normal discontinuities are only caused by the occlusion boundaries, which indicate that detecting the non-zero δ can be an effective measure to find discontinuities. However, this condition is too restrict for a real case where the normal map is obtained by PS or from a set of noisy point cloud. The computational error introduced by PS and the scanning error in the point cloud disturb the normal orientations, and produce the non-zero δ on the $\mathcal{P}_{i,j}$ which is expected to be continuous. It is obvious that the detection method based on a single ADD feature cannot distinguish discontinuity from noise as shown in Fig. 5 (Right). To improve the detection efficiency, a normal orientation feature is introduced in the next section.

3.3. Normal Angle Difference Feature

In general, surface discontinuities are mainly caused by the occlusion boundaries, where the ADD features are non-zero, and there exists the sharp orientation changes. Thus, both the non-zero δ and sharp orientation change are jointly considered to detect the discontinuity. To better quantify the change of orientations, a new feature called *normal angle difference* (NAD) is defined as the difference of normal orientations within a certain range, as given in Eq. (3):

$$\begin{aligned} \varphi_{i,j}^{\ell} &= \frac{\|\theta_{i,j}^{\ell}\|_{\infty}}{\|\theta_{i,j}^{\ell}\|_{-\infty}}, \\ &if\ \delta_{i,j}^{\ell} > 0,\ and\ \|\theta_{i,j}^{\ell}\|_{\infty} \equiv \theta_{i,j}^{\ell} \end{aligned}, \tag{3}$$

where $\varphi_{i,j}^{\ell}$ is the NAD value along ℓ in $\mathcal{P}_{i,j}$, and $\theta_{i,j}^{\ell}$ is,

$$\theta_{i,j}^{\ell} = [arccos(\mathbf{n}_l, \mathbf{n}_{ll}),\ arccos(\mathbf{n}_l, \mathbf{n}_r),\ arccos(\mathbf{n}_r, \mathbf{n}_{rr})]^T, \tag{4}$$

where $\mathbf{n}_{ll}$ and $\mathbf{n}_{rr}$ are the left-adjacent and right-adjacent normals of $\mathbf{n}_l$ and $\mathbf{n}_r$ respectively as shown in Fig. 3 (Right). $\| \cdot \|_{\pm\infty}$ is the $\pm\infty$-norm, which can be considered as the maximum/minimum operation.

3.4. Overall Detection Algorithm

The discontinuity detection that comprehensively considers both the ADD feature and the NAD feature can be implemented as:

Step 1) If $\delta_{i,j}^{\ell}$ is smaller than a given threshold τ_δ, then it is needed to check whether $\varphi_{i,j}^{\ell}$ exists or not. Otherwise, $\mathcal{P}_{i,j}$ is discontinuous along the direction $\boldsymbol{\ell}$.

Step 2) If $\varphi_{i,j}^{\ell}$ exists, it is needed to check the relationship between $\varphi_{i,j}^{\ell}$ and another threshold τ_φ. Otherwise, $\mathcal{P}_{i,j}$ is continuous along the direction $\boldsymbol{\ell}$.

Step 3) If $\varphi_{i,j}^{\ell}$ is bigger than τ_φ, then $\mathcal{P}_{i,j}$ is discontinuous along the direction $\boldsymbol{\ell}$. Otherwise, $\mathcal{P}_{i,j}$ is continuous along the direction $\boldsymbol{\ell}$.

Step 4) After all four directions are detected, the final splitting pattern of $\mathcal{P}_{i,j}$ is determined as follows: If two directions are detected as discontinuities for a $\mathcal{P}_{i,j}$, then the splitting pattern is shown in the first group of Fig. 2; Similarly, if three and four directions are detected as discontinuities for a $\mathcal{P}_{i,j}$, the associated splitting patterns are shown in the second and third group of Fig. 2, respectively.

Since $\mathbf{v}_{i,j}$ can be split according to the discontinuity detection condition, $\mathbf{z}(f_{i,j})$ in Eq. (1) is updated as $\mathbf{z}(f_{i,j}) = [z_{i,j}^{k_1},\ z_{i,j+1}^{k_2},\ z_{i+1,j+1}^{k_3},\ z_{i,j+1}^{k_4}]^T$. The superscript k represents the kth depth value along the z-axis.

4. Discontinuity Preservation

This section details the proposed DGP-based discontinuity preservation scheme for surface reconstruction from normals. A new light-weight energy model jointly considers the discontinuity detection parameters $\langle \tau_\delta, \tau_\varphi \rangle$ and the

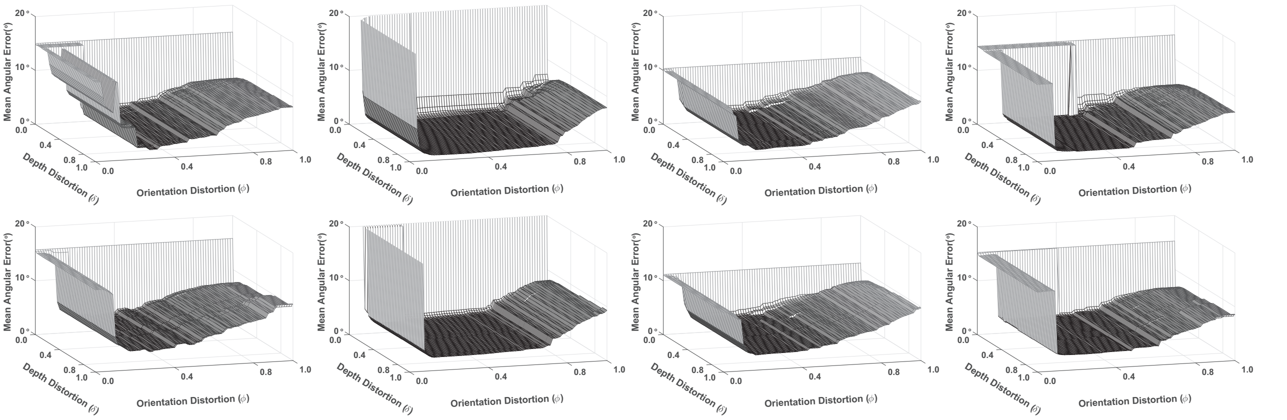

Bunny in Fig.11 | Face in Fig.12 | Hand in Fig.14 | Kitty in Fig.15

Figure 6. Energy surfaces of the *Kitty*, *Bunny*, *Hand*, and *Face* models. (Top): Energy surfaces of the original normal maps; (Bottom): Energy surfaces of the noisy normal maps with the variance of the Gaussian noise 3°.

overall surface reconstruction error. We also discuss how to efficiently compute the optimal detection parameters based on the prior knowledge.

4.1. Surface Reconstruction Optimization

In practice, it is infeasible to specify a different $\langle \tau_{\delta}, \tau_{\varphi} \rangle$ for each $\mathcal{P}_{i,j}$. For instance, considering that for a normal map with resolution of $W \times H$, the total number of unknown detection parameters is $2 \times (W \times H)$. To compute these unknown parameters, it greatly complicates the reconstruction process. Thus, we proposed to uniformly apply $\langle \tau_{\delta}, \tau_{\varphi} \rangle$ to the whole normal field, and the reconstruction energy model with respect to τ_{δ} and τ_{φ} can be formulated as,

$$argmin_{\{z_{i,j}^{k}\}} E(M(\tau_{\delta}, \tau_{\varphi})) \;\; s.t. \; \mathbf{n}(f_{i,j}) \equiv \mathbf{n}_{i,j}, \tag{5}$$

where $M(\tau_{\delta}, \tau_{\varphi})$ is the reconstructed mesh surface by Eq. (1), with the input normals whose discontinuities have been detected and meshed by τ_{δ} and τ_{φ} (*e.g.*, see Section 3.4). $E(M)$ is a function measuring the shape variation, and $\mathbf{n}(\cdot)$ returns the normal vector of a facet.

Fig. 6 shows the energy surfaces of four discontinuous models with the original normals and the noisy normals, respectively. It can be seen that a large $\langle \tau_{\delta}, \tau_{\varphi} \rangle$ results in a big reconstruction error. The reason is that when there are no discontinuities detected under the large detection parameters, $M(\cdot)$ is taken to be continuous. Strictly speaking, Xie *et al.* [24] is a special case of this research with infinite $\langle \tau_{\delta}, \tau_{\varphi} \rangle$. In addition, the result of a small $\langle \tau_{\delta}, \tau_{\varphi} \rangle$ is another extreme case where every $\mathcal{P}_{i,j}$ is determined to be discontinuous, and every $f_{i,j}$ is separated. In this case, the total number of unknown depth values is larger than the total number of normals, and Eq. (5) has no feasible solution. Consequently, the reasonable $\langle \tau_{\delta}, \tau_{\varphi} \rangle$ candidates are located in the middle range as shown in Fig. 6.

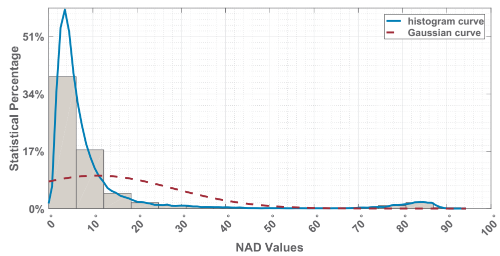

Figure 7. Histogram statistics of the ADD feature for the *Bunny* model in Fig. 5, where x axis is the relative value with respective to the range of ground-truth depth.

4.2. Determination of the Feasible Solution Space

The conventional way to compute the optimal $\langle \tau_{\delta}^{*}, \tau_{\varphi}^{*} \rangle$ is to use Newton's method for all possible values of δ and φ. However, it is time-consuming and easy to be stuck in a local minimum in the practical reconstruction, especially for the noisy normal map. Considering that a given surface to be reconstructed commonly has more continuities than discontinuities, the scope of searching $\langle \tau_{\delta}^{*}, \tau_{\varphi}^{*} \rangle$ can be shrunk based on the prior statistical knowledge of δ and φ.

The histogram statistics of the ADD feature $\delta = \{\delta_{i,j}^{\ell}\}$, which is computed by Eq. (2). An example of the histogram of δ for the *Bunny* model is given in Fig. 7. It can be seen that the ADD values are mainly concentrated in a narrow range with the mean value close to zero, and the rest are randomly distributed in a wide range with the upper limit approaching to the largest depth of the ground-truth. Supposing the ADD value with the largest probability is denoted as δ_0 (*i.e.* the $x-$ coordinate of the first bar in Fig. 7). A reasonable τ_{δ} is expected to be capable of filtering every $\mathcal{P}_{i,j}$ of δ_0 to avoid being disconnected. As a result, τ_{δ} can only be the value bigger than δ_0.

To specify the mathematical expression of the range of

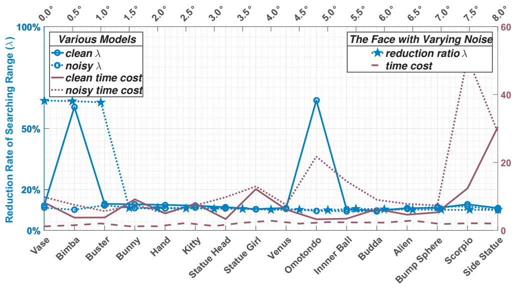

Figure 8. Average time cost and the reduction rate of searching range of $\langle \tau_\delta^*, \tau_\varphi^* \rangle$ for all 16 models, and for the *Face* model containing the Gaussian noise with variance ranging from $0°$ to $8°$.

τ_δ, we then fit the histogram h_δ by a Gaussian function g_δ. If the peaks of h_δ and g_δ are close, the optimal τ_δ^* should be larger than the mean value μ_δ of g_δ. If the peaks of h_δ and g_δ are far away from each other, it implies that the normal map contains heavy noise, which causes the distribution of δ diverging. Thus, in this research, the search range of the optimal τ_δ^* starts from $\mu_\delta + \frac{d_\delta}{2}$, where d_δ is the variance of g_δ. Eq. (6) gives the search range of τ_δ^*,

$$\Omega_\delta = \begin{cases} (\mu_\delta, \ \|\delta\|_\infty), & h_\delta(\mu_\delta) \approx max(h_\delta) \\ (\mu_\delta + \frac{d_\delta}{2}, \ \|\delta\|_\infty), & otherwise \end{cases} . \quad (6)$$

The distribution of the NAD feature is similar to that of the ADD feature, *i.e.*, mainly concentrated near small φ. Thus, the search range of the optimal τ_φ^* can be derived in a similar way. The difference is that we do not make analysis on all values of $\varphi = \{\varphi_{i,j}^\ell\}$ to compute the histogram h_φ. As we know that the patch, $\mathcal{P}_{i,j}$ having the largest probability δ_0, must be continuous, and the associated NAD feature $\varphi|_{\delta_0}$ rarely contains the optimal τ_φ^*. That is, the optimal τ_φ^* is impossible to be φ with the largest probability over the range of $\varphi|_{\delta_0}$. Eq. (7) gives the mathematical formulation of the search range of τ_φ^*,

$$\Omega_\varphi = \begin{cases} (\mu_\varphi, \ \|\varphi|_{\delta_0}\|_\infty), & h_\varphi(\mu_\varphi) \approx max(h_\varphi) \\ (\mu_\varphi + \frac{d_\varphi}{2}, \ \|\varphi|_{\delta_0}\|_\infty), & otherwise \end{cases} , \quad (7)$$

where h_φ is the histogram of $\varphi|_{\delta_0}$, μ_φ and d_φ are the mean and variance of the Gaussian function g_φ, which is fitted from h_φ.

To illustrate the effectiveness of Ω_δ and Ω_φ, we define the *range reduction ratio* as $\lambda = \frac{|\Omega_\delta|}{\|\delta\|_\infty} \times \frac{|\Omega_\varphi|}{\|\varphi\|_\infty}$, and test the proposed scheme on the *Face* model containing different noisy levels, and under the normals of 16 different models. The optimal detection parameters τ_δ^* and τ_φ^* are efficiently computed by the Monte Carlo method. As shown in Fig. 8, the searching scope can be shrunk to about 15% of the original one, even for the *Face* normals containing Gaussian noise of variance up to $8°$. The corresponding reconstruction error is provided in Fig. 9 and Fig. 10, where the mean angular error is close to $0°$.

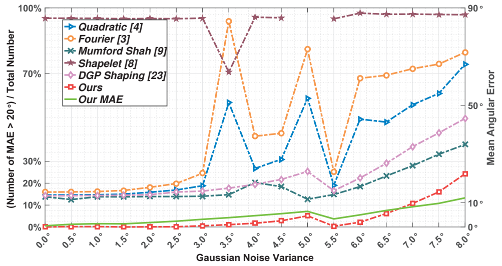

Figure 9. Reconstruction error rate of the *Face* model containing the Gaussian noise with variance ranging from $0°$ to $8°$. The indicating percentage is the ratio of the total number of reconstructed normals with angular error greater than $20°$ to the total number of input normals.

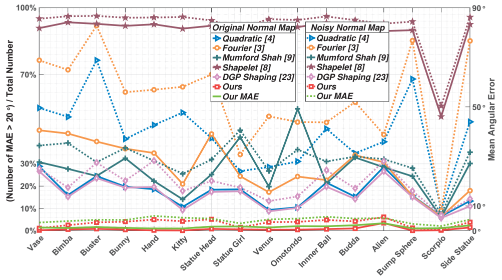

Figure 10. Average reconstruction error rate of the original/noisy normal maps of all 16 models. Results from the original normal maps are plotted by solid lines, while the noisy counterparts are denoted by dot lines.

5. Experimental Results

We implement the proposed approach in MATLAB under Intel i7 CPU with 3.40GHz and 8GB RAM, and compare the performance with five existing methods as mentioned in [13], including Quadratic [4], Fourier [3], Mumford Shah [9], Shapelete [8], and DGP shaping [24]. All the test examples are downloaded from Aim@Shape, and TurboSquid. For each 3D model, we first sample the normal vectors, render them via OpenGL, and then read back the associated depth map and the sampled normal map. To introduce the discontinuous artifacts into the tested examples, we simply introduce $[0, \ 0, \ 1]^T$ as surface normals in the regions where orientations are missing. For comparison, the input normal map serves as the ground-truth, where the angular errors between the input normals and the reconstructed ones are computed to measure the quality of reconstruction. Experimental results show that our approach is efficient and outperforms the other five methods. Detailed comparison and analysis are discussed as below.

Firstly, our approach is robust not only for the original normal maps, but also for the noisy ones. As shown in Fig. 9, the *Face* model containing the Gaussian noise with different variances ranging from $0°$ to $8°$ is used to test the

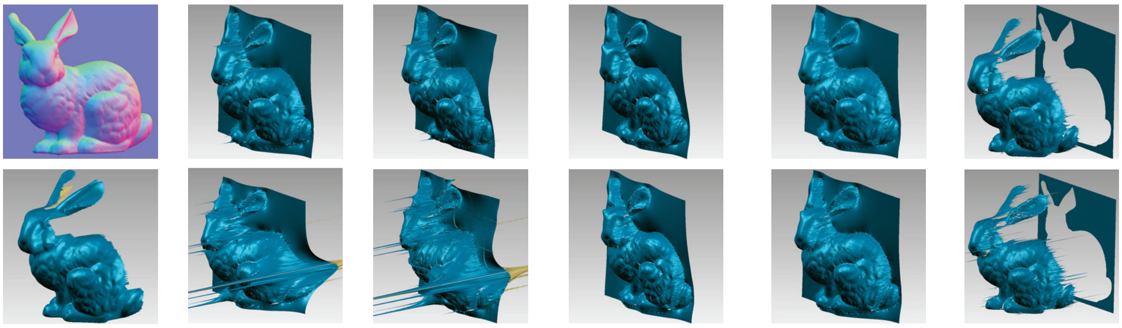

Ground-truth Quadratic [4] Fourier [3] Mumford Shah [9] DGP shaping [24] Ours

Figure 11. The *Bunny* model reconstructed by five different methods. (Top): Results from original normal map; (Bottom): Results from the normal map containing the Gaussian noise with variance of 4°.

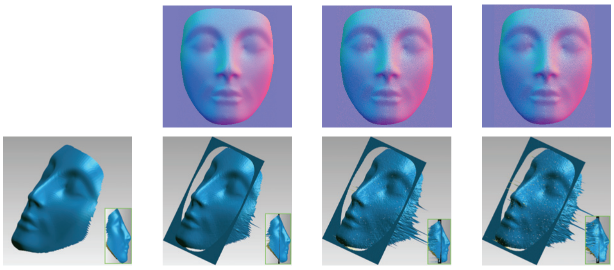

Ground-truth Var = 2° Var = 4° Var = 6°

Figure 12. Reconstructed surfaces of *Face* with varying Gaussian variances in Fig. 9. (Top): The input noisy normal maps; (Bottom): The reconstructed surfaces by our approach, where the associate side views are framed in a green box at the bottom-right corner.

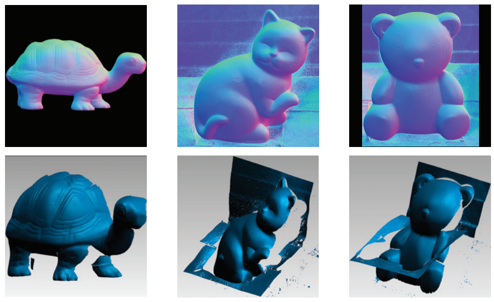

Turtle Bear Kitty

Figure 13. Reconstruction results of three real normal maps, i.e. the *Turtle* from the Harvard dataset, the *Cat* and *Bear* from the DiLiGenT dataset.

robustness of our method[2]. It can be seen that when the variance is increasing but less than 6°, the reconstruction error rate and mean angular error (MAE) [24] are kept below 5% and 5°, respectively. When the variance increases to 8°, the error rate and MAE are still below 10% and 10° respectively, which are much smaller than the other five methods. The related reconstructed surfaces are rendered in Fig. 12. It is clearly shown that all noisy cases of *Face* are satisfactorily separated from the background and accurately reconstructed.

Secondly, we also validate our approach on the *Bunny* and the *Hand* models as shown in Fig. 11 and Fig. 14, respectively. Each example is conducted on two types of inputs, the original normal map and the noisy one. It is obvious that all the other compared methods fail to break the discontinuous patches, thereby distort the reconstructed surface. We also observe that some methods are very sensitive to noise. For example, the reconstructed surfaces by Shapelete [8] fail to be rendered in the comparison as the errors are too large (see Fig. 9). In addition, we show more comparisons on 16 different models in Fig. 10. Obviously, the proposed method achieves the best performance. For the original normal maps, the reconstruction error rate and MAE are close to 0% and 0°, respectively. For the noisy counterparts, the reconstruction error rate and MAE are below 5% and 5°. All the other reconstruction results of our method are rendered and shown in Fig. 15.

Thirdly, we qualitatively evaluate our approach on three real-world PS normal models, including *Turtle* from the *Harvard* dataset, *Cat* and *Bear* from the *DiLiGenT* dataset, as shown in Fig. 13. The reconstructed mean angular errors are 2.16°, 4.13°, and 6.01°, respectively. It is worth noting that it is usually to remove the dark background to guarantee the reconstruction quality of computing normal by the traditional methods. Here, we still use the noisy (non-Gaussian) and unreliable background in the reconstructions of *Cat* and *Bear*.

Last but not the least, the proposed scheme can be efficiently solved. The reason is that Eq. (6) and Eq. (7) define an effective feasible solution space of $\langle \tau_\delta^*, \tau_\varphi^* \rangle$ for different normals. Especially, for a noisy normal map, the solution space is shrinking along with the increasing Gaussian variance. It means our approach can guarantee the reconstruc-

[2] Note that the maximum noise value can reach 30° with the Gaussian variance of 8°.

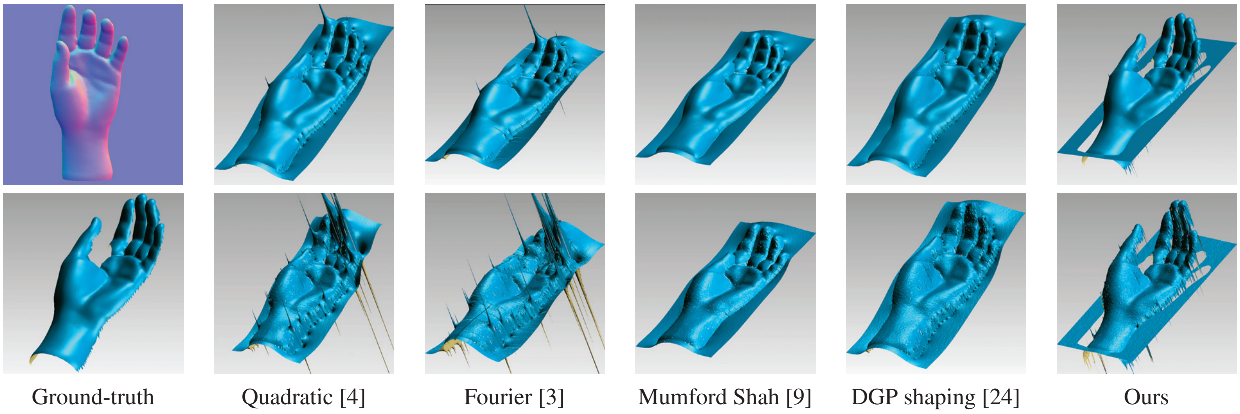

Figure 14. The *Hand* model reconstructed by five different methods. (Top): Results from original normal map; (Bottom): Results from normal map containing the Gaussian noise with variance of 6°.

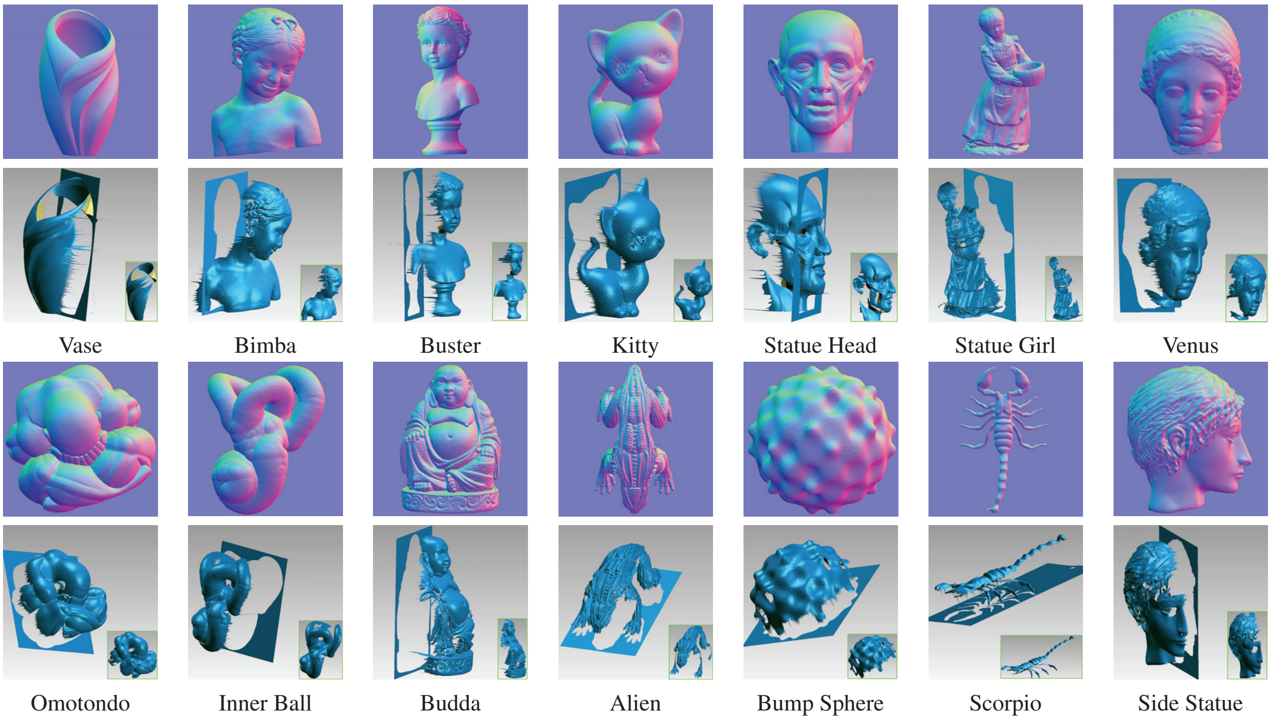

Figure 15. Reconstructed surfaces of different models by the proposed method, where the ground-truth under the same view are framed in a green box at the bottom-right corner.

tion efficiency for the noisy inputs without adding overhead complexity, which is validated in Fig. 8. The reconstruction time cost and the reduction rate of the searching range for the *Face* model with the different amount of noise are also plotted. It is observed that the time cost keeps the similar trend for different noises.

6. Conclusion

In this paper, we present an approach to address the problem of discontinuity preservation based on a discrete framework in surface reconstruction from normals. To determine the appropriate splittings of a discrete mesh, we introduce two new normal incompatibility features and design an efficient discontinuity detection scheme. In addition, we model the discontinuity preservation as a light-weight energy optimization problem, and shrink the feasible solution space to reduce the computation complexity based on the prior knowledge. Extensive experimental results show that the proposed method reconstructs various discontinuous surfaces efficiently and robustly.

References

[1] A. Agrawal, R. Raskar, and R. Chellappa. What is the range of surface reconstructions from a gradient field? In *ECCV 2006*, pages 578–591, 2006.

[2] A. Blake and A. Zisserman. *Visual reconstruction*. 1987.

[3] R. Framkot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *TPAMI*, (4):439–451, 1988.

[4] B. K. Horn and M. J. Brooks. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing*, 33(2):174–208, 1986.

[5] J.-W. Hsieh, H.-Y. M. Liao, M.-T. Ko, and K.-C. Fan. Wavelet-based shape from shading. *Graphical Models and Image Processing*, 57(4):343–362, 1995.

[6] B. Karacali and W. Snyder. Reconstructing discontinuous surfaces from a given gradient field using partial integrability. *Computer Vision and Image Understanding*, 92(1):78–111, 2003.

[7] B. Karacali and W. E. Snyder. Partial integrability in surface reconstruction from a given gradient field. In *Image Processing. 2002. Proceedings. 2002 International Conference on*, volume 2, pages II–II, 2002.

[8] P. Kovesi. Shapelets correlated with surface normals produce surfaces. In *ICCV 2005*, volume 2, pages 994–1001, 2005.

[9] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on pure and applied mathematics*, 42(5):577–685, 1989.

[10] H.-S. Ng, T.-P. Wu, and C.-K. Tang. Surface-from-gradients without discrete integrability enforcement: A gaussian kernel approach. *TPAMI 2010*, 32(11):2085–2099, 2010.

[11] N. Petrovic, I. Cohen, B. J. Frey, R. Koetter, and T. S. Huang. Enforcing integrability for surface reconstruction algorithms using belief propagation in graphical models. In *CVPR 2001*, volume 1, pages I–I, 2001.

[12] Y. Quéau and J.-D. Durou. Edge-preserving integration of a normal field: Weighted least-squares, tv and l^1 approaches. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 576–588, 2015.

[13] Y. Quéau, J.-D. Durou, and J.-F. Aujol. Variational Methods for Normal Integration. *Journal of Mathematical Imaging and Vision*, 60(4):609–632, 2018.

[14] R. F. Saracchini, J. Stolfi, H. C. Leitão, G. A. Atkinson, and M. L. Smith. A robust multi-scale integration method to obtain the depth from gradient maps. *Computer Vision and Image Understanding*, 116(8):882–895, 2012.

[15] Y. Sato, T. Okabe, and I. Sato. Bispectral photometric stereo based on fluorescence. In *CVPR 2012*, pages 270–277, 2012.

[16] B. Shi, K. Inose, Y. Matsushita, P. Tan, S.-K. Yeung, and K. Ikeuchi. Photometric stereo using internet images. In *3DV*, pages 361–368, 2014.

[17] T. Simchony, R. Chellappa, and M. Shao. Direct analytical methods for solving poisson equations in computer vision problems. *TPAMI*, 12(5):435–446, 1990.

[18] D. Terzopoulos. The computation of visible-surface representations. *TPAMI*, 10(4):417–438, 1988.

[19] Y. Wang, J. Bu, N. Li, M. Song, and P. Tan. Detecting discontinuities for surface reconstruction. In *ICPR 2012*, pages 2108–2111, 2012.

[20] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):191139, 1980.

[21] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *Asian Conference on Computer Vision*, pages 703–717. Springer, 2010.

[22] T.-P. Wu, J. Sun, C.-K. Tang, and H.-Y. Shum. Interactive normal reconstruction from a single image. *TOG 2008*, 27(5):119, 2008.

[23] T.-P. Wu and C.-K. Tang. Visible surface reconstruction from normals with discontinuity consideration. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 1793–1800, 2006.

[24] W. Xie, Y. Zhang, C. C. Wang, and R. C.-K. Chung. Surface-from-gradients: An approach based on discrete geometry processing. In *CVPR 2014*, pages 2195–2202, 2014.