

The following publication W. Xie, M. Wang, X. Qi and L. Zhang, "3D Surface Detail Enhancement from a Single Normal Map," 2017 IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 2017, pp. 2344-2352 is available at <https://doi.org/10.1109/ICCV.2017.255>.

3D Surface Detail Enhancement from A Single Normal Map

Wuyuan Xie^{1,*}, Miaohui Wang^{2,*}, Xianbiao Qi¹, Lei Zhang¹

¹Dept. of Computing, The Hong Kong Polytechnic University, Hong Kong, China

²College of Information Engineering, Shenzhen University, Guangdong, China

wuyuan.xie@gmail.com, mhwang@szu.edu.cn, qixianbiao@gmail.com, cslzhang@comp.polyu.edu.hk

Abstract

In 3D reconstruction, the obtained surface details are mainly limited to the visual sensor due to sampling and quantization in the digitalization process. How to get a fine-grained 3D surface with low-cost is still a challenging obstacle in terms of experience, equipment and easy-to-obtain. This work introduces a novel framework for enhancing surfaces reconstructed from normal map, where the assumptions on hardware (e.g., photometric stereo setup) and reflection model (e.g., Lambertian reflection) are not necessarily needed. We propose to use a new measure, angle profile, to infer the hidden micro-structure from existing surfaces. In addition, the inferred results are further improved in the domain of discrete geometry processing (DGP) which is able to achieve a stable surface structure under a selectable enhancement setting. Extensive simulation results show that the proposed method obtains significantly improvements over uniform sharpening method in terms of both subjective visual assessment and objective quality metric.

1. Introduction

In 3D reconstruction, the surface details are mainly limited to visual sensors of digital equipment. Specifically, micro-structure of surface is degraded by sampling and quantization in the digitalization process. Consequently, although many dense-based 3D reconstruction techniques (i.e., photometric stereo [5, 16] and structured light [11]) are used to perform the reconstruction on a pixel-by-pixel basis, the reconstructed surface quality is still affected by the input data, not to mention other sparse-based 3D reconstruction algorithms (i.e., depth camera [8], stereo/multi-view stereo vision [12, 14]). However, as the trend of 3D reconstruction in industry is more inclined to use low-cost technology, how to improve the surface quality from existing depth data is a challenging but meaningful research topic.

Unlike shape information, the texture pattern of 3D object cannot be directly represented by the depth values, but can be reflected by the distribution of the pixel gray-scale values observed on the object appearance. How to obtain these gray-scale values from the depth map needs to employ the reflection model which defines the relationship between the observed pixel value and the normal vector under a given lighting condition [24]. Reflection model is the backbone of the most-recent works on the surface enhancement, and these works can be divided into two categories according to their applications: 1) shade rendering [10, 17, 23] in computer graphics, and 2) *photometric stereo* (PS) [21, 22] based surface texture enhancement. In shade rendering, a complex reflection model is usually employed, where the captured pixel is rendered by adjusting the lighting direction. Note that these methods do not really change the surface geometry structure, but make varying lighting conditions such that the observed object appearance seems to have richer details.

PS-based enhancement methods utilize one camera to capture a set of images under different lighting conditions, and the high-resolution normal map can be reconstructed from the raw data according to the *Lambertian* reflection [24]. However, PS-based methods change the surface geometry structure in a live capture way. Tan *et al.* [22] proposed to improve the surface micro-structure by photometric stereo, where the lighting direction within one pixel is estimated by the Gaussian mixture model (GMM). Li *et al.* [15] proposed to refine the surface by multi-view stereo where triangular meshing of surface is a necessary condition. Authors in [18, 13, 1, 9] used the normal map to repair the scanned sparse depth map as well as to recover the surface details. Significantly, all these methods rely on the related hardware which will cause two limitations: 1) it cannot enhance 3D surface from existing depth data; 2) it is strongly dependent on a reflection model. As a result, they are not satisfied for general 3D surface enhancement, which is expected to perform on existing depth or normal map without taking assumption on any reflection models.

In this paper, we present an efficient surface enhance-

* Authors contributed equally to this work.

This work is supported by China NSFC grant (no. 61672446).




Figure 1: Enhanced 3D surface results: the left is a 3D surface of the *Scolar* normal map from Harvard dataset [26]; the right is our directly enhanced result of the same resolution as inputs.

ment framework which can adaptively select parameters to amplify the current normal angle based on its texture density. We summarize our main contributions as three aspects: 1) the proposed method is independent of hardware, reflection model, or training dataset. Given a normal map as input, our method generates a 3D surface with enhanced appearance; 2) we propose the concept of *angle profile* to measure the local distribution of micro surface structure, and deformation of the profile can effectively change the quality of surface details; 3) we introduce the local-global strategy based on *discrete geometry processing* (DGP) [3, 25] to solve the difficulty of one-angle-to-many-profiles which is triggered by simultaneously deforming all angle profiles. Simulation results show that our method can greatly improve the fine-grained appearance of 3D surface reconstruction in terms of structural similarity index measure (SSIM).

2. Surface Details in Normal

From a microscopic point of view, surface details presented as depth relationship within a small appearance patch is more easily perceived in normal field. We embrace this observation and propose to convert the input as below if it is depth data,

$$\mathbf{N}(i, j) = \left[\frac{\partial \mathbf{D}(i, j)}{\partial x}, \frac{\partial \mathbf{D}(i, j)}{\partial y}, -1 \right]^T, \quad (1)$$

where \mathbf{D} and \mathbf{N} are depth map and normal map respectively, and (i, j) is the pixel coordinate.

2.1. Normal Angle

The surface normals can effectively describe the shape as well as surface details. There is also evidence [7] that the shape and surface details are determined by the low-frequency and high-frequency components of \mathbf{N} . In general,



Figure 2: Amplification of normal angle.

the low-frequency normal \mathbf{N}_L can be computed by the low-pass filtering on \mathbf{N} , and the high-frequency normal can be considered as \mathbf{N} itself. It is worth noting that when adjusting \mathbf{N} and fixing \mathbf{N}_L at the same time, we can change the surface details without deforming the surface shape.

In this section, we measure the relationship between \mathbf{N} and \mathbf{N}_L by the angle between them,

$$\Theta(i, j) = \arccos(\mathbf{N}(i, j) \cdot \mathbf{N}_L(i, j)), \quad (2)$$

where $\Theta(i, j)$ is normal angle at (i, j) . It is observed that when Θ increases, surface structure is more visually appealing. Hence, a direct way to enhance the surface is to enlarge Θ with respect to a constant \mathbf{N}_L . Fig. 2 illustrates the normal angle amplification in terms of Θ , \mathbf{N} , and \mathbf{N}_L , where normal angle reflects the contrast of surface structure. Normal map can be updated as \mathbf{N}^* under a new Θ^* ,

$$\mathbf{N}^*(i, j) = \mathbf{N}(i, j) + \frac{\mathbf{N}(i, j) - \mathbf{N}_L(i, j)}{|\mathbf{N}(i, j) - \mathbf{N}_L(i, j)|} \sqrt{2(1 - \cos(\Theta^*(i, j)))}. \quad (3)$$

2.2. Angle Profile

According to Eq. (3), quality of the enhanced surface depends on the selection of Θ^* . Intuitively, *uniform sharp-*




Figure 3: Illustration of *angle profiles*: the top row is the *angle profiles* with *profile sharpness* φ ; the bottom row is the *sharpened angle profiles* highlighted in red. $\varphi = \arccot \left| \frac{a}{a-b} \right|$, where a and b are two constants defined in Eq. (6).

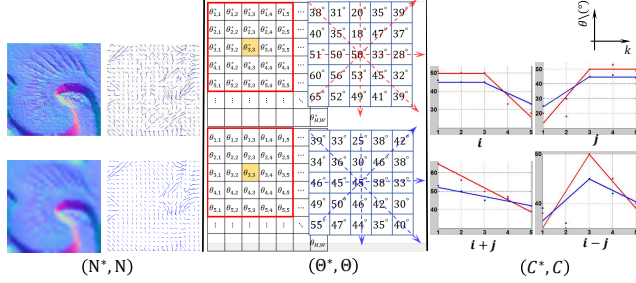


Figure 4: Illustration of fitting *angle profile* along four directions, e.g., \mathbf{i} , \mathbf{j} , $\mathbf{i} + \mathbf{j}$, and $\mathbf{i} - \mathbf{j}$, on a 5×5 patch centering on $(3, 3)$: the top row is high-quality normal map \mathbf{N}^* with *angle profiles* \mathbf{C}^* highlighted in red solid line; the bottom row is the input normal map \mathbf{N} with *angle profiles* highlighted in blue solid line.

ening is setting $\Theta^* = k \times \Theta$, with the scaling factor satisfying $k > 1$. Such an updating method is not an effective way to improve surface details, especially when the surface has uneven structure. Simulation results in Section 4 (i.e., Fig. 9, $\Theta^* = 1.5 \times \Theta$) show the unsatisfying performance in this case. The main reason is that the updating of Θ^* should be adaptive according to the local density of surface structure. Consequently, we propose to use *angle profile* to measure the local surface structure as the normal angle distribution along certain direction on a small patch. We observe that sharpening the *angle profile* appropriately to adjust the existing normal angle is of use to generate high-quality surface details. We demonstrate how to extract *angle profile* from a given normal map as follows.

Suppose we have computed the normal angle Θ from the normal map \mathbf{N} by Eq. (2). Now considering a $m \times m$ angle patch centering on (i, j) , we reorganize the angle values along direction \mathbf{d} as a $m \times 1$ vector,

$$\theta_{i,j}^{\mathbf{d}} = \{\Theta(s, t) \mid s = \tan(\mathbf{d} \cdot \mathbf{i})(t - j) + i\}, \quad (4)$$

or,

$$\theta_{i,j}^{\mathbf{d}} = [\theta_1, \dots, \theta_m]^T,$$

where $\theta_{i,j}^{\mathbf{d}}$ is the *angle vector*, \mathbf{d} is a normalized vector with four possible directions, i.e., \mathbf{i} , \mathbf{j} , $\mathbf{i} + \mathbf{j}$, or $\mathbf{i} - \mathbf{j}$ (\mathbf{i} and \mathbf{j} are unit vectors parallel to \mathbf{x} and \mathbf{y} axis), and t ranges

from $j - \frac{m-1}{2}$ to $j + \frac{m+1}{2}$. Since we are interested in the distribution of $\theta_{i,j}^{\mathbf{d}}$, we consider $\theta_{i,j}^{\mathbf{d}}$ as a mapping operation from element index to its value: $\theta_k = \theta_{i,j}^{\mathbf{d}}(k)$ with $k \in [1, \dots, m]$. Apparently, such a mapping lacks generalization ability to describe the local distribution of angle values. Extensive experimental results show that by applying a linear regression model f on $\theta_{i,j}^{\mathbf{d}}$, we can well simulate the mapping curve as one of the nine contours illustrated in Fig. 3.

$$\mathbf{C}_{i,j}^{\mathbf{d}} = f(\theta_{i,j}^{\mathbf{d}}), \quad (5)$$

where $\mathbf{C}_{i,j}^{\mathbf{d}}$ is named as *angle profile* centering on (i, j) , and f is a piecewise linear function with a stationary point on $\langle k, \theta_k \mid k = \frac{m+1}{2} \rangle$.

$$f(\theta) = \min_{\{a,b,d\}} \sum_k \|ak + b\theta_k - d\|^2. \quad (6)$$

Eq. (6) is applied for the left and right half of $\theta_{i,j}^{\mathbf{d}}$ to fit the target *angle profile*, respectively. An example of fitting *angle profile* along four directions is illustrated in Fig. 4.

Adjusting the normal angle by deforming $\mathbf{C}_{i,j}^{\mathbf{d}}$ will affect surface details, and we hence propose to quantize the *profile sharpness* as the angle φ as shown in Fig. 3. For simplicity, we denote all the *angle profiles* from \mathbf{N} as a set, \mathbf{C} (i.e., $\mathbf{C}(i, j | \mathbf{d}) = \mathbf{C}_{i,j}^{\mathbf{d}}$), and the associated *profile sharpness* as Φ (i.e., $\Phi(i, j | \mathbf{d}) = \varphi_{i,j}^{\mathbf{d}}$), respectively.

We observe that, if normal map \mathbf{N} has a high-quality update \mathbf{N}^* , then on the same pixel coordinate and the same direction, the fitted *angle profiles* have similar contours but different sharpness, where *angle profile* of \mathbf{N}^* is more sharper. This is because, the more clearer surface details, the bigger variance of normal angles. Conversely, if *angle profile* appears to be horizontal (Fig. 3 (i)), it means that the surface is originally smooth and should be preserved after enhancement. Therefore, we arrive at the conclusion that if non-horizontal \mathbf{C} is appropriately sharpened, the resulting normals can be used to produce a fine-grained 3D surface, which is also confirmed by simulation results in Section 4.

2.3. Sharpening Weight

$\mathbf{C}_{i,j}^{\mathbf{d}}$ can be sharpened by decreasing its *profile sharpness* $\varphi_{i,j}^{\mathbf{d}}$, and the decreasing measure is defined as *sharpening*




Figure 5: Sharpening for angle profile.

weight $w_{i,j}^d$, i.e. $w_{i,j}^d \cdot \varphi_{i,j}^d \leq \varphi_{i,j}^d$. Given $w_{i,j}^d$, the sharpening operation is required to satisfy the following two constraints to avoid over-sharpening on surface: 1) the area of the right triangle should change slowly (i.e., the shaded triangle in Fig. 5); 2) only translation is allowed on the horizontal angle profile. We implement such a deformation on C^1 according to its angle profile type as shown in Fig. 3, which can be divided into four cases:

Case 1: C is type (i). It indicates that there is no variance on the associated normal angles, and the original surface should be smooth. So sharpening operation is not performed.

Case 2: C is type (a) or type (b) with profile sharpness φ and sharpening weight w . We observe that if the sharpening operation causes instant change on the triangle area of C , it may spoil the original surface structure. To avoid this negative effect, we first initialize C as an equal-area triangle \tilde{C} with profile sharpness $w \times \varphi$ (i.e., blue triangle in Fig. 5 left). Meanwhile, if C is directly changed to \tilde{C} , it may result in over-smoothing on surface. So we build a smooth transition between v_p and v_e by introducing a buffer tail $\psi = \widehat{v_p v_e}$ (i.e., solid green line in Fig. 5), which is defined as an arc with one tangent parallel to the k -axis. Notice that v_p is the cross point between C and \tilde{C} , thus it is easy to derive the circle form of ψ . Consequently, the updated C is,

$$C^*(k) = \begin{cases} \tilde{C}(k), & k \geq k_p \\ \psi(k), & k < k_p \end{cases} \quad (7)$$

Case 3: C is type (c) or type (d). The updated C^* is obtained by performing the sharpening operation on the left half and right half of C independently.

Case 4: C belongs to type (e) \sim type (h) in Fig. 3. C^* is computed by first performing sharpening operation on the part with profile sharpness φ , and conjoin the rest.

The sharpened angle profiles are illustrated on the bottom row of Fig. 3. An enhanced normal map N^* requires

¹In this section, the sharpening operation is discussed on a local patch only, and for simplicity, the subscripts of (i, j) and d are ignored without further statement.

each angle profile being sharpened appropriately, which means that the selection of $w_{i,j}^d$ is a crucial factor. In Section 3, we describe how to obtain an optimal weight map $W : W(i, j | d) = w_{i,j}^d$ according to a user-selectable sharpening scale.

3. Refinement as Discrete Geometry Processing

We propose to refine a 3D surface to a user-selectable level λ . Different from *uniform sharpening*, we expect the enhanced surface should keep similar texture pattern as inputs. The problem can be formulated as,

$$\min_{\{\Theta(i,j)\}} \mathbf{E}(\mathbf{N}) \quad s.t. \quad W(i, j | d) = \frac{\max(\Phi)}{\lambda \times \Phi(i, j | d)}, \quad (8)$$

where $\mathbf{E}(\mathbf{N})$ is the cost function measuring the variation of surface structure given a normal map \mathbf{N} , Θ is computed by Eq. (2), and Φ and W are the associated profile sharpness map and sharpening weight map, respectively. We observe that the statistical distribution of Φ can effectively describe a whole surface structure. Given λ , a reasonable W can be guaranteed by $\max(\Phi)/\Phi$.

Since Φ is dependent on Θ , we take the iterative strategy to solve Eq. (8). The current Θ is computed based on the previous Φ , where the iteration process is terminated until Φ reaches stable values. In addition, it is easy to fall into a local minimum to solve Eq. (8) using the traditional optimization method. Thus, we propose to find the optimum Θ in discrete geometry processing (DGP) domain.

3.1. DGP Setup

For an input \mathbf{N} and a user-given setting λ , suppose the associated angle profile C have been sharpened as C^* under an initialized W . Then, we can compute a new Θ^* which locally satisfies each angle profile in C^* . Apparently, a proper Θ^* cannot be obtained by directly mapping the angle values from C to C^* , because C holds the one-to-many relationship between angle and angle profile, where a normal angle is constrained by multiple angle profiles (i.e., from both neighboring patches and different directions). Fortunately, following the local/global formulation [19, 20], the non-linear constrained geometry processing problems can be solved by iteratively applying the local projection and the global blending step [3]. Our solution is partly inspired by this local/global strategy, where the normal angle is adjusted according to the involved sharpened angle profiles in Eq. (7), and the profile sharpness map Φ is redistributed according to a new angle map Θ^* based on the global least-square optimization. Θ^* directly affects the constraint in Eq. (8), and the updated constraint will trigger a new round of iteration.




Figure 6: Local projection and global blending under four constraints.

3.2. Local/global Solution

We restate our input-output as: input is a 2D normal angle map Θ and a user-selectable scale λ ; output is an enhanced angle map Θ^* . To solve this problem, we generate a five-step DGP solver as:

- Step 1 Compute \mathbf{C} according to Θ by Eq. (5), and initialize the sharpening weights as $\mathbf{W}(i, j | d) = \frac{\max(\Phi)}{\lambda \times \Phi(i, j | d)}$, where Φ is calculated from \mathbf{C} in Subsection 2.2.
- Step 2 Use \mathbf{W} to sharpen the *angle profile* as \mathbf{C}^* (see Subsection 2.3).
- Step 3 Obtain a new Θ^* according to \mathbf{C}^* by the local/global strategy (see Subsection 3.3). For a single normal angle θ in Θ , the local projection step is performed by mapping values from \mathbf{C} to \mathbf{C}^* . Suppose there is a total number of g *angle profiles* associated with θ , then θ will have total g different target values. As illustrated in Fig. 6, if the related \mathbf{C}^* is considered as a feasible region, the local projection step actually projects the current θ onto its closest point. Thereafter, the global blending step is applied to determine a compromised position of θ according to all different projected values. All updated θ contributes to a new Θ^* ;
- Step 4 Use Θ^* to update the normal map as \mathbf{N}^* (see Eq. (3)), where \mathbf{N}_L is extracted from the original input and kept as a constant during each iteration.
- Step 5 Set Θ as Θ^* , and return to *Step 1* to conduct the next round of optimization.

After iteratively performing *Step 1* to *Step 5*, the DGP solver will converge to an optimal \mathbf{W}^* which determines a stable surface structure under the enhancement level λ . Then, the associated output Θ^* is used to achieve an enhanced \mathbf{N}^* according to Eq. (3). The main reason is that the local/global strategy can obtain an optimal solution for the problem in Eq. (8), where the local step avoids the problem of finding a “good” initial guess that is needed by most optimization methods in solving non-convex problems, but

easily stuck at a local optimum when given a “bad” initialization. Detailed implementation of *Step 3* is given in Subsection 3.3.

3.3. Formulation in DGP

In the global blending of *Step 3* in Subsection 3.2, we aim to solve an improved Θ^* where each angle vector can be fitted into a curve with the same *profile sharpness* as \mathbf{C}^* . Specifically, a profile \mathbf{C} corresponds to a certain θ . Accordingly, the sharpened \mathbf{C}^* has the counterpart θ^* . For a given Θ^* , each pair of $\theta_{i,j}^d$ and $\theta_{i,j}^{*d}$ is expected to exhibit the same sharpness of *angle profiles*. Consequently, a straightforward formulation is used to minimize the following cost function,

$$\mathcal{L}(\{\theta_{r,c}^l\}) = \sum_d \sum_{i,j} \|\theta_{i,j}^d - \theta_{i,j}^{*d}\|^2, \quad (9)$$

By minimizing the above function, we can obtain a normal angle map Θ^* with the associated *angle profiles* close to the target \mathbf{C}^* . It is worth noting that directly solving Eq. (9) will make the DGP solver converge slowly, as $\theta_{i,j}^d = \theta_{i,j}^{*d}$ is a hard constraint condition. The convergence speed can be improved by translating $\theta_{i,j}^d$ and $\theta_{i,j}^{*d}$ into the local coordinate with the origin on their own center. So, the corresponding reformulation of Eq. (9) is,

$$\mathcal{L}(\{\theta_{r,c}^l\}) = \sum_d \sum_{i,j} \|(I - \frac{1}{m}\mathbf{1})(\theta_{i,j}^d - \theta_{i,j}^{*d})\|^2, \quad (10)$$

where m is the length of θ , I is a $m \times m$ unit matrix, and $\mathbf{1}$ is a $m \times m$ matrix with all elements equal to 1.

A more efficient way to solve Eq. (10) is developed by reformulating it into the matrix form, $\mathcal{L}(\{\theta_{r,c}^l\}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$, where \mathbf{A} is a matrix derived from $(I - \frac{1}{m}\mathbf{1})\theta_{i,j}^d$, \mathbf{b} is a vector derived from $(I - \frac{1}{m}\mathbf{1})\theta_{i,j}^{*d}$, and \mathbf{x} is the vector containing all unknown angle values of Θ^* .

4. Experimental results

We implement our algorithm in *MATLAB R2014a* and evaluate its performance on real depth or normal data from scanning or *photometric stereo*. All the simulation results are obtained on a uniform platform, *Intel TM2 CPU with 3.16GHz and 8GB RAM*. We compare our approach with *uniform sharpening* by setting $\Theta^* = 1.5 \times \Theta$ in Eq. (3). The detailed experiments are discussed as below.

Fig. 7 illustrates the general performance in enhancing five real examples, and the quantitative results are tabulated in Table 2. Since SSIM has been generally considered as a better perceptual quality metric for 2D natural image or depth image [2, 4] than PSNR, we compute the SSIM score for comparison based on the depth map (*i.e.*, all compared pair depths are normalized into the same range). For each surface patch, we compute the *profile sharpness* of inputs




Figure 7: SSIM scores along *profile sharpness* of 5 models listed in Table 2. Blue, red and green lines correspond to $\frac{1}{2} \times \frac{1}{2}$, $\frac{1}{3} \times \frac{1}{3}$, and $\frac{1}{4} \times \frac{1}{4}$ downsampled inputs, respectively. The lines labelled with “+” and “•” denote our method and *uniform sharpening*. Notice that, our enhancement from $\frac{1}{3} \times \frac{1}{3}$ input achieves the same quality as from $\frac{1}{2} \times \frac{1}{2}$, where the associated lines are nearly overlapping.

and the mean SSIM of the results, and plot them into a “sharpness-SSIM” curve. We use blue, red, and green lines to denote three downsampling ratios $\frac{1}{2} \times \frac{1}{2}$, $\frac{1}{3} \times \frac{1}{3}$, and $\frac{1}{4} \times \frac{1}{4}$, respectively. We use symbols “+” and “•” to distinguish our results from the *uniform sharpening*. As shown in Fig. 7, our method shows great stability in $\frac{1}{2} \times \frac{1}{2}$ and $\frac{1}{3} \times \frac{1}{3}$ downsampling, with high SSIM score for different *profile sharpness*. However, for a lower downsampling (i.e., $\frac{1}{4} \times \frac{1}{4}$), our method fails to guarantee the quality of enhancement, especially for the patch with high profile sharpness. Consequently, we conclude that performance of our algorithm is dependent on the variance of the *angle profile* inferred from the input normal data. To illustrate this relation, we also plot the sharpness curves of the *vase* surface under three different downsampling ratios (i.e., $\frac{1}{2} \times \frac{1}{2}$, $\frac{1}{3} \times \frac{1}{3}$, and $\frac{1}{4} \times \frac{1}{4}$) as shown in Fig. 8. For comparison convenience, all the downsampled normal maps are interpolated to the same resolution as the original one. To visually compare contour of the curves, we re-index the patch in the ascending order based on the *profile sharpness* from the original normal map. As illustrated in Fig. 8, the trends of the sharpness curve of $\frac{1}{2} \times \frac{1}{2}$ and $\frac{1}{3} \times \frac{1}{3}$ downsampling are similar to the original, although they are more flatter, while the curve of $\frac{1}{4} \times \frac{1}{4}$ downsampling is almost to be a straight line.

We apply our method for 11 real surfaces including 9 scanned depth data downloaded from *Aim@Shape* (see Fig. 9, Fig. 10, and Fig. 11), and 2 normal data by *photometric stereo* (see the *Lincoln cent* and *Cloth* in Fig. 11). Since our algorithm only accepts the normal map as the input, for the depth data, we use *MeshLab* [6] to re-mesh it and convert it into normal. For each example as shown in Fig. 9, we first downsample the original normal map and enhance it to the same resolution as the original one. For visualization




Figure 8: *Profile sharpness* comparisons among the original, $\frac{1}{2} \times \frac{1}{2}$, $\frac{1}{3} \times \frac{1}{3}$, and $\frac{1}{4} \times \frac{1}{4}$ downsampled normal maps of the *vase* model in Fig. 9.

Table 1: Mean SSIM scores of depth maps in Fig. 9.

	$\frac{1}{2} \times \frac{1}{2}$	$\frac{1}{3} \times \frac{1}{3}$	$\frac{1}{4} \times \frac{1}{4}$
<i>uniform's</i>	0.8384	0.9542	0.5943
<i>Ours</i>	0.9898	0.9914	0.6172

Table 2: Mean SSIM scores of depth maps in Fig. 10[†].

	<i>Vase</i>	<i>Eros</i>	<i>Wall_1</i>	<i>Life_mask</i>	<i>Rame.</i>
<i>uniform's</i>	0.9542	0.9864	0.9538	0.9816	0.9837
<i>Ours</i>	0.9914	0.9993	0.9884	0.9991	0.9885

[†] All the compared enhancements are based on $\frac{1}{3} \times \frac{1}{3}$ downsampled inputs.

purpose, we use the reconstruction method [25] to estimate the 3D surface from normal. We set the normal angle patch as 5×5 for fitting the *angle profile*. Larger size of patch (i.e. $> 7 \times 7$) is not suggested, because some hidden structure is lost during *angle profile* fitting. As shown in Fig. 9, for $\frac{1}{2} \times \frac{1}{2}$ downsampling, surface details are recovered (see the close-up view in the red box in Fig. 9), where the SSIM score is up to 0.9898 in comparison with the original data (see Tabel 1). For the case of $\frac{1}{3} \times \frac{1}{3}$ downsampling, the visual quality of enhancement is almost as good as that of $\frac{1}{2} \times \frac{1}{2}$, although the input loses more information. For the case of $\frac{1}{4} \times \frac{1}{4}$ downsampling, our result in terms of SSIM drops to 0.6172, but still gives better result than *uniform sharpening*. This is because under $\frac{1}{4} \times \frac{1}{4}$ downsampling, the reserved hidden structure information is not enough to obtain a proper *angle profile*.

Fig. 10 shows the comparison results of $\frac{1}{3} \times \frac{1}{3}$ downsampling for the models of *Eros*, *Wall_1*, *Life_mask*, and *Rame*. Our method is conducted on each model of the top row, and the bottom row gives the corresponding enhance-




Figure 9: Comparison of the *vase* surface enhancements under different downsampling ratios. All partial close-up views are in red box. The most left column is the ground-truth with 298×298 resolution with side view on its top. The second to the last column are enhancements on $\frac{1}{2} \times \frac{1}{2}$, $\frac{1}{3} \times \frac{1}{3}$, $\frac{1}{4} \times \frac{1}{4}$ down-sampled normals, where the first row is input, the second row is obtained by *uniform sharpening*, and the third row is obtained by our method.




Figure 10: The 3D surface enhancements for $\frac{1}{3} \times \frac{1}{3}$ downsampling of the model *Eros*, *Wall_1*, *Life_mask*, and *Rame*. The top row is the downsampled inputs with resolution of 100×100 , 102×102 , 100×100 , and 35×35 , respectively. The bottom row is the enhanced results by our method.

ments. We can see that substantial hidden structures are visually improved by our method. The related SSIM scores




Figure 11: Surface enhancement comparisons. There are six original normal maps from the right-most column to the left-most are *dragon scales*, *hair of Eros*, *fish tail*, *part of human lips*, *Lincoln cent*, and *cloth*, respectively. The first four models are depth data scanned from plaster models, and the last two models are the normal data obtained by *photometric stereo*. The first row is the original inputs, the second row is results by *uniform sharpening* with $\Theta^* = 1.5 \times \Theta$, and the third row is our method with $\lambda = 0.3$.




Figure 12: Comparison of applying image sharpening on normal map. The normal map is taken as an image and applied three up-sampling (3×3) methods bicubic, bilinear, nearest followed by image sharpening. The associated depth SSIM values are indicated below the images.

are tabulated in Table. 2.

We also evaluate our method on the original surface without downsampling. Fig. 1 indicates our enhancement results on the publicly available normal map of *scolar* from Harvard dataset [26]. Fig. 11 shows the visual comparison between our method and *uniform sharpening*. There is no SSIM score for the quantitative analysis, because the related ground-truth is not available. However, it can be observed that our enhancements show more high-frequent information without over-sharpening. Note that, for *Lincoln Cent* and the *cloth*, their original normal maps are obtained by *photometric stereo*.

We conduct the experiemnts by taking the normal map as an image and applying up-sampling followed by image sharpening. Fig. 12 shows the comparison of reconstruction surface on the Eros model under three up-sampling methods *bicubic*, *bilinear*, and *nearest*, where the upsampling scale

is 3×3 . The associated depth SSIM values are about 0.9813, 0.9797, 0.9788 and 0.9993 respectively.

5. Conclusion

This paper presents a new method to enhance 3D surface details for general given normal maps. We address the fine-grained 3D enhancement by introducing *angle profile* which makes our method independent of photometric stereo as well as the reflection model assumption. The hidden surface structure is obtained from the raw data based on its *angle profile*. Meanwhile, we propose to use the DGP method to further refine the enhanced surface. Extensive simulation results show that our method greatly improves the fine-grained details of the raw depth.

References

- [1] K. Achuta, T. Vage, S. Boxin, and R. Ramesh. Polarized 3d: High-quality depth sensing with polarization cues. In *ICCV*, 2015.
- [2] A. Benoit, P. Le Callet, P. Campisi, and R. Cousseau. Quality assessment of stereoscopic images. *EURASIP journal on image and video processing*, 2008(1):1–13, 2009.
- [3] S. Bouaziz, M. Deuss, Y. Schwartzburg, T. Weise, and M. Pauly. Shape-up: Shaping discrete geometry with projections. In *Computer Graphics Forum*, volume 31, pages 1657–1667, 2012.
- [4] M.-J. Chen, C.-C. Su, D.-K. Kwon, L. K. Cormack, and A. C. Bovik. Full-reference quality assessment of stereopairs accounting for rivalry. *Signal Processing: Image Communication*, 28(9):1143–1155, 2013.
- [5] D. Cho, Y. Matsushita, Y.-W. Tai, and I. Kweon. Photometric stereo under non-uniform light intensities and exposures. In *ECCV*, 2016.
- [6] P. Cignoni, M. Corsini, and G. Ranzuglia. Meshlab: an open-source 3d mesh processing system. *Ercim news*, 73(45-46):6, 2008.
- [7] P. Cignoni, R. Scopigno, and M. Tarini. A simple normal enhancement technique for interactive non-photorealistic renderings. *Computers & Graphics*, 29(1):125–133, 2005.
- [8] Y. Cui, S. Schuon, S. Thrun, D. Stricker, and C. Theobalt. Algorithms for 3d shape scanning with a depth camera. *TPAMI*, 35(5):1039–1050, 2013.
- [9] N. Diego, R. Szymon, D. James, and R. Ravi. Efficiently combining positions and normals for precise 3D geometry. *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH 2005)*, 2005.
- [10] R. Fattal, M. Agrawala, and S. Rusinkiewicz. Multiscale shape and detail enhancement from multi-light image collections. *ACM Transactions on Graphics (TOG)*, 26(3):51, 2007.
- [11] M. Gupta, A. Agrawal, A. Veeraraghavan, and S. G. Narasimhan. Structured light 3d scanning in the presence of global illumination. In *CVPR*, 2011.
- [12] X. Hu and P. Mordohai. A quantitative evaluation of confidence measures for stereo vision. *TPAMI*, 34(11):2121–2133, 2012.
- [13] M. Kiyoshi and A. Yoshimitsu. Depth image enhancement using local tangent plane approximations. In *CVPR*, 2015.
- [14] A. Kuhn, H. Hirschmüller, D. Scharstein, and H. Mayer. A tv prior for high-quality scalable multi-view stereo reconstruction. *IJCV*, pages 1–16, 2016.
- [15] S. Li, S. Y. Siu, T. Fang, and L. Quan. Efficient multi-view surface refinement with adaptive resolution control. In *ECCV*, 2016.
- [16] F. Lu, Y. Matsushita, I. Sato, T. Okabe, and Y. Sato. From intensity profile to surface normal: photometric stereo for unknown light sources and isotropic reflectances. *TPAMI*, 37(10):1999–2012, 2015.
- [17] S. Rusinkiewicz, M. Burns, and D. DeCarlo. Exaggerated shading for depicting shape and detail. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 1199–1205, 2006.
- [18] B. Sema, F. Xinyi, and R. Szymon. Merge2-3D: Combining multiple normal maps with 3D surfaces. *International Conference on 3D Vision*, 2014.
- [19] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. *ACM Transactions on Graphics (TOG)*, 23(3):399–405, 2004.
- [20] R. W. Sumner, M. Zwicker, C. Gotsman, and J. Popović. Mesh-based inverse kinematics. *ACM transactions on graphics (TOG)*, 24(3):488–495, 2005.
- [21] P. Tan, S. Lin, and L. Quan. Resolution-enhanced photometric stereo. In *ECCV*, 2006.
- [22] P. Tan, S. Lin, and L. Quan. Subpixel photometric stereo. *TPAMI*, 30(8):1460–1471, 2008.
- [23] R. Vergne, R. Pacanowski, P. Barla, X. Granier, and C. Shlick. Improving shape depiction under arbitrary rendering. *IEEE Transactions on visualization and computer graphics*, 17(8):1071–1081, 2011.
- [24] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):191139–191139, 1980.
- [25] 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.
- [26] X. Ying, C. Ayan, B. Ronen, G. Steven, J., J. David, W., and Z. Todd. Photometric stereo dataset. <http://vision.seas.harvard.edu/qsfs/Data.html>.