

Detailed Surface Geometry and Albedo Recovery from RGB-D Video Under Natural Illumination

Xinxin Zuo^{1,2,*}Sen Wang^{2,1,*}Jiangbin Zheng²Ruigang Yang^{1,3}¹ University of Kentucky
Lexington, KY, US² Northwestern Polytechnical University
Xi'an, P.R.China³ Baidu
Beijing, P.R.China

{xinxin.zuo, sen.wang}@uky.edu

zhengjb@nwpu.edu.cn

ryang@cs.uky.edu

Abstract

In this paper we present a novel approach for depth map enhancement from an RGB-D video sequence. The basic idea is to exploit the photometric information in the color sequence. Instead of making any assumption about surface albedo or controlled object motion and lighting, we use the lighting variations introduced by casual object movement. We are effectively calculating photometric stereo from a moving object under natural illuminations. The key technical challenge is to establish correspondences over the entire image set. We therefore develop a lighting insensitive robust pixel matching technique that out-performs optical flow method in presence of lighting variations. In addition we present an expectation-maximization framework to recover the surface normal and albedo simultaneously, without any regularization term. We have validated our method on both synthetic and real datasets to show its superior performance on both surface details recovery and intrinsic decomposition.

1. Introduction

The availability of affordable depth sensors has sparked a revolution in many areas of computer vision, such as human computer interactions, robotics, and video analysis. Among these 3D modeling has probably received the most benefits from this advancement of sensors. Nevertheless, the current generation of depth sensors still suffers from limited resolution and accuracy. As a result, fine-scale structural details of an object cannot be recovered.

Many researchers have exploited high quality color image as guidance to enhance the depth map, such as denoise and up-sampling [28, 11, 40, 6, 25]. Among these approaches, the fusion of depth maps with shading information contained in color images has shown its effectiveness in recovering surface geometry details. Most of these

methods implement shading refinement on single RGB-D frame on the basis of Shape-from-Shading (SfS) techniques [41, 27, 37]. However, the inherent ambiguity for SfS still exists and smooth regularization term is usually incorporated. Another problem is that an albedo image is needed to predict the appearance, which is also unknown. Typically a constant albedo assumption is made. In order to handle varying albedos, previous shading-based refinement methods often approached this chicken-and-egg problem by assuming prior assumptions or by enforcing particular albedo regularizers. However, these regularizers are heuristic and may not work all the time. In this paper we propose to take advantage of an RGB-D sequence to uniquely resolve the normal maps and albedo images.

We only make two basic assumptions about the object, first it is primarily diffuse and second it is acquired by casual movement of the object (not the camera). In this way, the object's movement induces illumination changes in the image sequence, which is critical to resolve the surface normal and albedo without any ambiguity. It resembles the photometric stereo. But instead of controlling the light when imaging the static object, we move the object under general natural lighting. This kind of cue has been exploited in multi-view photometric stereo [10] and shape from video [22, 32, 42]. However, they have the environmental lighting constrained to be calibrated directional light and the object is experiencing turntable motion or the motion is already calibrated. On the contrary our approach works under natural lighting with the object under arbitrary motion, which makes our method more widely used in everyday environment.

Given the captured RGB-D sequence, first we try to align the RGB-D sequence and find the correspondences among the images using a novel robust matching technique. Then the environmental lighting is estimated using the intensity ratios of the aligned sequence, which effectively factors out the impact of varying albedo. Finally, we formulate an Expectation-Maximization framework in which the surface normal and its albedo can be calculated robustly, in the pres-

*The first two authors contributed equally to this work

ence of some non-Lambertian reflection or cast shadow. A detailed surface mesh is obtained after integration of the initial depth map with estimated normal map.

The main contribution is that we utilize the dynamic photometric information along the sequence to recover the surface details beyond the resolution of current depth sensors. Compared to previous depth enhancement schemes that use the color information, our method, to the best of our knowledge, is the least restrictive. It allows arbitrary surface albedo, does not require controlled or calibrated lighting or turntable capture. To achieve these, we make two technical contributions. The first is a novel image registration scheme that is robust to lighting variations and the second is an EM optimization scheme to produce per-pixel normal and albedo map under general lighting.

2. Related Work

In this section, we will review the previous works in two related topics: surface geometry or depth enhancement with shading information, and intrinsic image decomposition.

2.1. Shape from shading and photometric stereo for surface or depth enhancement

The Shape-from-shading (SfS) problem has long been studied since the pioneering work by Horn [16]. It aims to estimate surface normal (and then indirectly surface shape) from a single image. There are various regularization terms or prior assumptions [1, 2] that have been enforced to deal with the inherently ill-posed problem.

Recent methods have shown that SfS can be used to refine the noisy depth map captured from RGB-D cameras [14, 41]. The inherent ambiguity of SfS is not resolved exactly, but with the initial depth close to the real surface, Wu [37] and Roy [27] have achieved good performance in recovering surface details. Varying albedo poses another challenge as it needs to be factored out before lighting estimation and shading refinement. Some [14] assumes uniform or piecewise constant albedo. Yu [41] deals with this by clustering a fixed set of discrete albedos before optimizing geometry. A better, yet more complex strategy, is to simultaneously optimize for unknown albedos and refine geometry [18]. There are also previous works that adopt the shading constraints to improve the coarse 3D shape reconstructed using multi-view stereo [36]. The major difference between our method and these prior works is that we fully exploit lighting effects contained in a video sequence to uniquely and simultaneously determine the surface albedo and normal in a pixel-wise manner. Therefore we are able to deal with arbitrary albedo.

Photometric stereo methods have been developed to compute the surface normal using multiple images of a scene taken under different or controlled illumination [38, 39]. Unlike SfS, photometric stereo is a well-defined

problem, which is then incorporated to enhance the raw depth map [15, 43] captured from the current depth sensor. Wu [35] and Zhou [46] have also shown its effectiveness under multi-view setup or for images acquired from the internet [31]. Besides, some have used the IR images instead of color images for normal estimation [34]. Recently Chatterjee [8] exploited the IR image and proposed a factorization method to handle objects with varying albedo. In these approaches, the objects are kept to be static or captured almost under the same viewpoint with different lighting conditions. Different from them we allow the object to rotate arbitrarily under uncontrolled environment.

2.2. Intrinsic decomposition

There are some works about intrinsic decomposition that focus on separating the albedo map from shading image, to which our work is also related. The problem of intrinsic image decomposition, first introduced by Barrow [4], is to separate reflectance and illumination from a single image. It is again an ill-posed problem since there are two unknowns for every observation. Additional constraints must be adopted to make the problem solvable. The Retinex theory [23] is widely used for this purpose. It assumes that shading variations are mostly low frequency while albedo changes are mostly high-frequency. Based on this assumption, many approaches have been proposed. For example, Tappen et al. [33] train a classifier from reflectance and shading data sets. Global priors, such as the reflectance sparsity, are developed [13, 5, 44] and they perform clustering or enforce non-local smoothness on image reflectance. These priors or regularization terms are not guaranteed to work well in all cases, especially when the albedo variation is significant (e.g., a textured surface).

More information has been used to reduce the ambiguity. For example, researchers propose to use RGB-D image as input and depth maps or normals are taken as additional cues [9, 17, 3]. Lee et al. [24] employed temporal albedo consistency constraints for RGB-D video. Another solution is to leverage multiple images of the scene taken under varying lighting conditions [19, 21, 20]. However, without any information about the surface geometry and environmental lighting, the problem is still ill-posed with regularization terms or pairwise constraints needed.

The main difference between our method and the prior works is that we can solve normal and albedo for each pixel independently, *without the need for additional regularization terms*.

3. Preliminary Theory

While environmental lighting can be arbitrarily complex, the appearance of a diffuse object can be described by a low dimensional model [29]. Under this assumption, the shading function s for Lambertian reflectance can be modeled

as a quadratic function of the surface normal with A , b , c represented as the lighting parameters.

$$s(\mathbf{n}) = \mathbf{n}^T A \mathbf{n} + b^T \mathbf{n} + c \quad (1)$$

Then the captured image is generated by multiplying the shading function with surface albedo $\rho(p)$

$$I(p) = \rho(p) \cdot s(\mathbf{n}(p)) \quad (2)$$

Given a single image as observation, it may not be feasible to recover the surface normal and albedo since for each pixel we have three equations with five unknowns to be estimated. Photometric stereo, with more lighting variations, is a typical solution to resolve the ambiguity. Mathematically, the surface normal and its albedo can be computed by minimizing the following objective function that is formulated for each pixel independently under natural lighting condition. No smoothness or albedo regularization is needed.

$$E(\mathbf{n}, \rho) = \sum_k \|\rho \cdot (\mathbf{n}^T A_k \mathbf{n} + b_k^T \mathbf{n} + c_k) - I_k\|^2 \quad (3)$$

The underlying principle of our enhancement method is based on photometric stereo, but we do not need to change the lighting condition with the object being static; instead we captured the RGB-D sequence of the object under arbitrary motion in uncalibrated natural illumination. Suppose we keep the first frame as the reference frame, and we can find the correspondences for pixels along the sequence. For example, for pixel p in the reference frame, its correspondence in frame k is $W(p)$. The appearance of the pixel $W(p)$ is generated as,

$$\begin{aligned} I_k(W(p)) &= \rho(p) \cdot ((R_k \mathbf{n})^T A (R_k \mathbf{n}) + b^T (R_k \mathbf{n}) + c) \\ &= \rho(p) \cdot (\mathbf{n}^T (R_k^T A R_k) \mathbf{n} + (b^T R_k) \mathbf{n} + c) \end{aligned} \quad (4)$$

where ρ is the albedo for pixel p which equals to pixel $W(p)$ and \mathbf{n} is surface normal under reference frame coordinate. R_k is the rotation from the reference frame to frame k .

Therefore the changes of lighting induced by the object motion provide valuable cues to recover the surface normal and its albedo resembling photometric stereo.

4. Pipeline

An overview of our depth enhancement and albedo recovery framework is shown in Figure 1.

First, we fuse every $M = 20$ RGB-D frames via Kinect-Fusion [26] to generate N key frame depth maps. They are smoother and more accurate than the raw depth maps. The extrinsic parameters between these key frames are computed and then refined after bundle adjustment. A robust

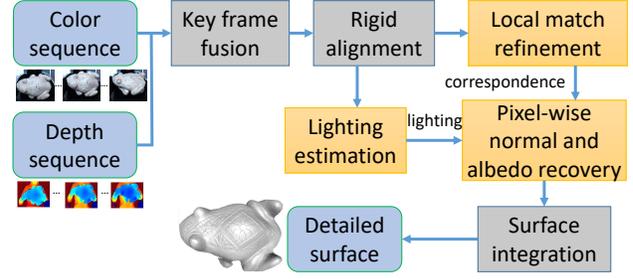


Figure 1: System Pipeline.

pixel matching strategy is proposed to deal with misalignment. For lighting estimation, the entire sequence is used to make the estimation more robust. Finally, given the computed lighting and correspondences along the sequence, we recover the surface normal and albedo image under our robust EM framework. The recovered normal can be integrated with the reference key frame depth map to generate surface model with much more structural details.

5. Approach

There are three major parts in our approach, including searching correspondences among the images, lighting estimation, and normal and albedo recovery. They will be described in details in the following sections.

5.1. Robust Pixel Matching

The key frame depth maps $D_0 \sim D_N$ are obtained via depth fusion and the corresponding color images are denoted as $I_0 \sim I_N$. We set the first frame as the reference frame ($D_{\text{ref}} = D_0$, $I_{\text{ref}} = I_0$). We need to establish correspondences between pixels in reference frame and those in other frames.

5.1.1 Rigid alignment

First, the global rigid transformation from this reference frame to other frames are calculated by detecting SIFT or ORB features followed by feature matching. These extrinsic parameters are further refined with bundle adjustment and finally we get the rotation $R_0 \sim R_N$ and translation matrix $T_0 \sim T_N$ from the reference frame to other frames.

5.1.2 Lighting insensitive robust match

These key frames can be warped into the reference frame given the current transformation. However misalignments still exist after bundle adjustment as shown in Fig. 2(c), which is caused by the imprecise depth maps and the imperfect synchronization between the captured color and depth sequence. Optical flow is often used as a solution to find correspondences between two images. Considering that the

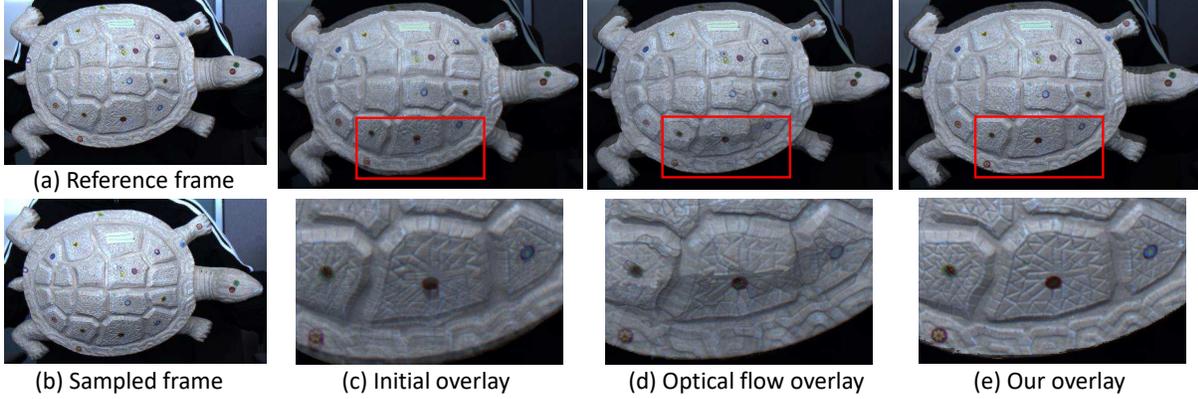


Figure 2: Demonstration of matching. (a) is the reference frame and (b) is one sampled key frame. Image(b) is warped to the reference frame with current transformation, and (c) displays the warped image overlaid with image(a). (d) shows the overlaid result using the flow map computed from warped image and the reference image. (e) is the overlaid image after applying our proposed lighting insensitive robust matching.

misalignment may be severe, we have tried to use large displacement optical flow [7] to find the correspondences between the warped image and the reference image. However, since the consistency assumption is not maintained in our case, the alignment has got even worse in some part with great illumination changes as displayed in Fig. 2(d). Therefore, we have developed a robust method to deal with these issues.

Suppose we have the reference depth map and color image denoted as D_{ref} and I_{ref} respectively. For each pixel $p = (u, v)$ in I_{ref} , its current corresponding pixel(q) after bundle adjustment in image I_k is computed as,

$$\lambda \begin{bmatrix} q \\ 1 \end{bmatrix} = K \left(R_k (K^{-1} * \begin{bmatrix} u \\ v \\ D_{\text{ref}}(u, v) \end{bmatrix}) + T_k \right) \quad (5)$$

The corresponding pixel p in I_{ref} and q in I_k may not be the correct correspondence because of the misalignment. Therefore, we implement a local search strategy to find its best matching pixel in I_k .

For each pixel p in I_{ref} , we set a searching region around it and find its best match in I_k via NCC. However, the intensity consistency is not preserved as the object is subject to arbitrary movements. This makes the original NCC not suitable for matching in this case. To deal with this problem, we apply chromaticity normalization in the color image to eliminate the effect of lighting variations [12] and use the normalized images for matching. For each pixel p , its appearance is generated as,

$$I_{\text{ch}}(p) = \rho_{\text{ch}}(p) * s(p) \quad \text{ch} \in \{R, G, B\}, \quad (6)$$

in which $s(p)$ is the shading function that accounts for the lighting or normal variation.

So the chromaticity normalization is implemented as,

$$I_{\text{ch}}^{\text{cn}}(p) = \frac{I_{\text{ch}}(p)}{I_R(p) + I_G(p) + I_B(p)} \quad \text{ch} \in \{R, G, B\}, \quad (7)$$

After this normalization, NCC can then be applied for matching insensitive to the photometric inconsistency induced by lighting factor.

Besides, the color image I_{ref} is warped to the color frame I_k under the guidance of D_{ref} and we get the warped color image I_{ref_k} . The NCC patch matching is implemented in I_{ref_k} with I_k instead of using I_{ref} directly. Since I_k and I_{ref_k} are in the same viewpoint, the fattening effect of NCC is successfully avoided.

Although for each pixel in I_{ref} (or I_{ref_k}) we can find the corresponding pixel in I_k that has the largest matching score, we cannot guarantee they are always the correct correspondence. To tackle this problem, we only keep the pixels that are reliable and use these pixels as control vertices to deform all the other pixels to find their correct correspondences.

Our criteria of reliable matches is that, 1) the largest matching score should be larger than thres_S ; 2) the difference between the largest score and second largest score of local peaks should be larger than thres_Δ . If these principles are maintained, the pixel in the searching region that has the largest score is chosen as the correspondence. thres_S is set to be 0.75 and thres_Δ is 0.05 in this paper for all the experiments.

Next we use these reliable matches as control vertices to deform the image I_{ref_k} so that it has an optimal match with I_k . As for each control vertices o_l in I_{ref_k} , the deformation

function is defined as

$$f(o_l) = o_l + \Delta_l, \quad (8)$$

where Δ_l is the motion vector between the optimal correspondence and its initial correspondence in I_k .

For other pixels the deformation is formulated via bilinear interpolation with control vertices [45],

$$f(u) = u + \sum_l (\theta_l^u \Delta_l) \quad (9)$$

The interpolation coefficients θ_l^u is set according to the distance to control vertices and only neighboring vertices will affect the deformation.

Finally, our objective function is defined to maintain the photo consistency of the two normalized images.

$$E(\Delta) = \sum_p \left(I_{\text{ref}_k}^{cn}(p) - I_k^{cn}(f(p, \Delta)) \right) + \lambda \sum_l \|\Delta_l - \hat{\Delta}_l\|^2 \quad (10)$$

For the control vertices $\hat{\Delta}$ is the initial deformation vector between the current optimal correspondence obtained from matching and its initial correspondence. λ is the control weight set to be 10 in this paper. Since we have good initials Δ , the optimization will converge quite fast.

The matching result can be seen in Fig. 2(e).

5.2. Lighting estimation

The unknown albedo poses challenges for lighting estimation. In this paper, we employ the aligned color sequence and depth maps for robust lighting estimation, eliminating the need to make assumptions about albedo. With the aligned color images we can compute the ratio images to the reference image, for which the albedo is eliminated. For each pixel p in I_{ref} , suppose its corresponding pixel in I_k is denoted as q , then the ratio value is computed as,

$$\frac{I_k(q)}{I_{\text{ref}}(p)} = \frac{\rho(q)(\mathbf{n}_q^T R^T A R \mathbf{n}_q + b^T R \mathbf{n}_q + c)}{\rho(p)(\mathbf{n}_p^T R^T A R \mathbf{n}_p + b^T R \mathbf{n}_p + c)} = \frac{\mathbf{n}_q^T R^T A R \mathbf{n}_q + b^T R \mathbf{n}_q + c}{\mathbf{n}_p^T R^T A R \mathbf{n}_p + b^T R \mathbf{n}_p + c} \quad (11)$$

Therefore, the environmental lighting can be achieved from the following minimization,

$$\arg \min_{A, b, c} \sum_k \sum_{p \in I_{\text{ref}}} \gamma_p \left\| \frac{\mathbf{n}_q^T R^T A R \mathbf{n}_q + b^T R \mathbf{n}_q + c}{\mathbf{n}_p^T R^T A R \mathbf{n}_p + b^T R \mathbf{n}_p + c} - \frac{I_k(q)}{I_{\text{ref}}(p)} \right\|^2 \quad (12)$$

The normal \mathbf{n} are approximated using normals computed with the key frame depth maps. The weighting term γ_p is set to prevent the effects of dark pixels which are noisy and might be caused by shadow.

5.3. Normal and albedo recovery

With the key frame color images all aligned into the reference frame ($I_0^W \sim I_N^W$), estimated environmental lighting (A, b, c), and object rotation matrix $R_0 \sim R_N$ for each

frame, we are ready to recover the surface normal and its albedo. For each pixel p in the frame if we denote observations as $O = \{I_k^W(p)\}$, then our goal is to find the optimal albedo ρ and normal \mathbf{n} confronting the pixel observations. We drop the index of pixel locations for simplicity in the following description. The objective function can be defined as:

$$E(\mathbf{n}, \rho | O) = \sum_k \|\rho \cdot s_k(\mathbf{n}) - I_k^W\|^2, \quad (13)$$

$$s_k(\mathbf{n}) = \mathbf{n}^T R_k^T A R_k \mathbf{n} + b^T R_k \mathbf{n} + c \quad (14)$$

The surface normal and albedo can be estimated after minimization of the above function. However, the outliers have not been taken into consideration. They will affect the result if the observations violate the Lambertian assumption or are in cast shadow. To deal with these outliers, we introduce a set of hidden states $H_k = \{0, 1\}$ indicating whether the observation is actually generated by the Lambertian model. An expectation-maximization (EM) algorithm is developed to solve the problem. While our formulation is inspired by [38], we extend it from its original directional light assumption to general lighting. More specifically, we denote the parameters to be estimated as $\Omega = \{\mathbf{n}, \rho, \sigma, \alpha\}$ and the observation probability conditioned on parameters Ω is given as,

$$P(O_k | \Omega) = \alpha \cdot \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\|\rho \cdot s_k(\mathbf{n}) - I_k^W\|^2}{2\sigma^2}\right) + (1 - \alpha) \cdot \frac{1}{C} \quad (15)$$

$P(H_k = 1) = \alpha$ is the prior probability of H_k indicating the proportion of observations generated by the Lambertian model. $\frac{1}{C}$ is the probability as being an outlier.

The posterior probability of the hidden variable H_k given parameters Ω' in current iteration and the observation O_k is computed in every E-step,

$$\begin{aligned} \omega_k &= P(H_k = 1 | O_k, \Omega') \\ &= \frac{\alpha \cdot \exp\left(-\frac{\|\rho \cdot s_k(\mathbf{n}) - I_k^W\|^2}{2\sigma^2}\right)}{\alpha \cdot \exp\left(-\frac{\|\rho \cdot s_k(\mathbf{n}) - I_k^W\|^2}{2\sigma^2}\right) + \frac{1 - \alpha}{C}} \end{aligned} \quad (16)$$

Next, in the following M-step, we maximize the complete-data log-likelihood given the marginal distribution H_k obtained from the E-step.

$$\begin{aligned} P(\Omega) &= \sum_k \log P(O_k, H_k = 1 | \Omega) \omega_k \\ &+ \sum_k \log P(O_k, H_k = 0 | \Omega) (1 - \omega_k) \\ &= \sum_k \log\left(\frac{\alpha}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\|\rho \cdot s_k(\mathbf{n}) - I_k^W\|^2}{2\sigma^2}\right)\right) \omega_k \\ &+ \sum_k \log\left(\frac{1 - \alpha}{C}\right) (1 - \omega_k) \end{aligned} \quad (17)$$

To maximize the above function, we set the first derivative of P with respect to α , σ and ρ equal to zero. In this way, the updating rules for these parameters are obtained,

$$\begin{aligned} \alpha &= \frac{1}{N} \sum_k \omega_k \\ \sigma &= \sqrt{\frac{\sum_k \|\rho \cdot s_k(\mathbf{n}) - I_k^W\|^2 \omega_k}{\sum_k \omega_k}} \\ \rho &= \frac{1}{\sum_k s_k(\mathbf{n})^2 \omega_k} \sum_k s_k(\mathbf{n}) \cdot \omega_k \cdot I_k^W \end{aligned} \quad (18)$$

Since the function P is nonlinear to surface normal \mathbf{n} , the updated normal is achieved by fixing other parameters and solving the following energy minimization.

$$\arg \min_{\mathbf{n}} \sum_k \|\rho \cdot (\mathbf{n}^T R_k^T A R_k \mathbf{n} + b^T R_k \mathbf{n} + c) - I_k^W\|^2 \omega_k \quad (19)$$

The above EM iterative optimization process is performed until no further improvement on the recovered normal and albedo. The initial parameter of α and σ is set to be 0.75 and 0.05 respectively for all the datasets used in this paper.

Finally, the normal is integrated with the reference depth map to get enhanced surface geometry with structural details [43].

5.4. Implementation details

As a preprocessing step, the object is first segmented from the image by integrating both color and depth information into GrabCut [30] framework. We manually masked the first frame with the rest of frames segmented automatically.

We implement most parts of our framework in Matlab and it takes us approximately 800s to process a dataset with 500~600 frames. Considering that the normal and albedo is computed in pixel-wise manner, the running time could be reduced further with parallel computation.

6. Experimental results

In the experiments, we validate our method on synthetic and real datasets with quantitative and qualitative comparison.

6.1. Synthetic data

In this section we use Stanford Bunny as our synthetic model and preform a quantitative comparison between our method and a shading refinement method [27]. First given the 3D model, twenty images together with the depth maps are rendered under natural illumination. The rendered ground truth depth maps are over smoothed to filter out the

structural details. Those smoothed depth maps and rendered color images are taken as input for our method.

Fig. 3 shows the comparison results on recovered normal map and surface. The first column is the reference color map and over smoothed mesh (displayed as normal map). These are the input for shading refinement method. The output of shading refinement method is displayed in the second column. The texture copy artifacts are caused by imperfect separation of albedo and shading layers. In comparison, the surface normal can be recovered successfully with pixel-wise recovery method with quite small error shown in the third column. The albedo map computed from our method together with its error map is demonstrated in the last column.

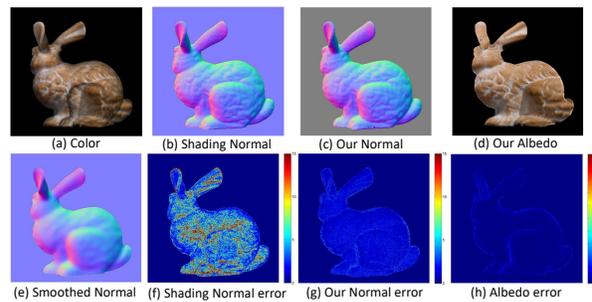


Figure 3: Results on synthetic Bunny model. (a) is the rendered color image of the reference frame; (e) shows the normal map of the ground truth mesh after over smoothing; (b) is the normal map computed after applying shading refinement on the reference frame with its error map displayed in (f); (c) and (g) are the normal map and its corresponding error map achieved by our method. Our recovered albedo map and its error map is also demonstrated in (d) and (h) respectively.

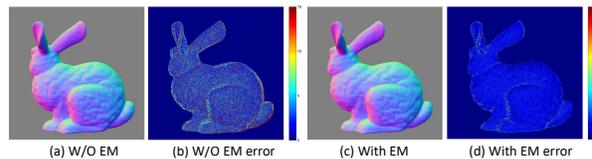


Figure 4: Results when adding salt and pepper noise. (a) shows the computed normal map without our EM framework and (b) is its error map; The normal and error map after applying our EM optimization are shown in (c) and (d) respectively.

Fig. 4 is shown to demonstrate the effectiveness of our EM framework for robust normal recovery in the presence of outliers. We have picked four out of those twenty images randomly and added the salt and pepper noise with 0.50 density. It means the abrupt noise will affect approximately fifty percent of the image pixels. As we can see from the first two columns, the recovered normal map without EM

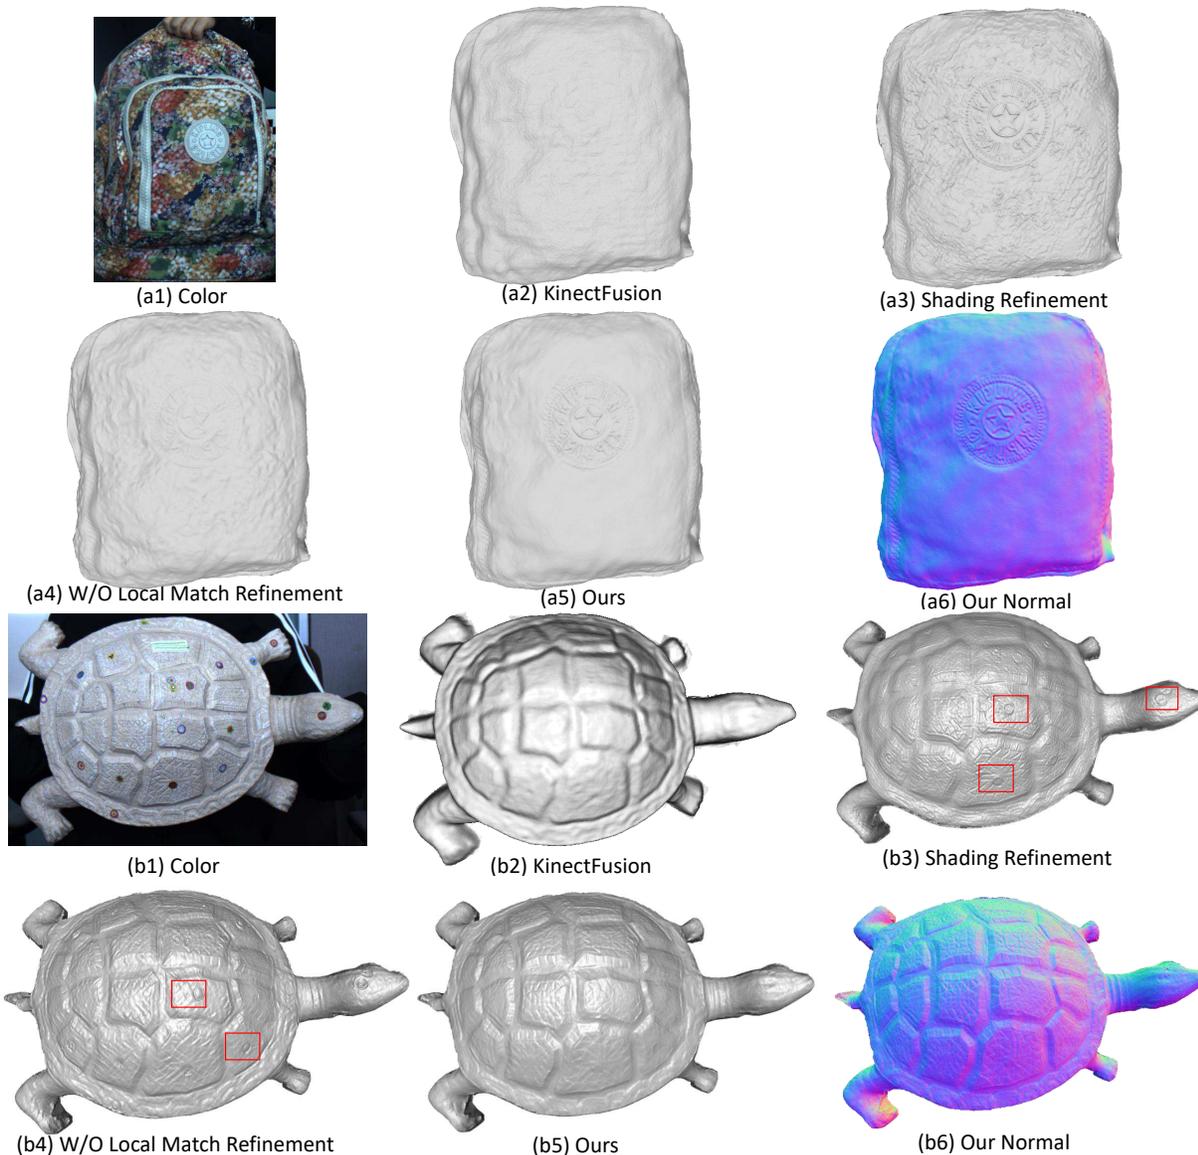


Figure 5: Comparison results on Backpack and Turtle model. (a1) and (b1) are the reference color images of Backpack and Turtle respectively. The output from KinectFusion is shown in (a2) and (b2). The results computed by shading refinement method are displayed in (a3) and (b3). (a4) and (b4) are the meshes acquired using our method but without applying our locally robust matching procedure. Finally, (a5) and (b5) are the meshes achieved by our approach. The normal map is given in (a6) and (b6).

optimization is noisy (the mean error is 8.37 degree), while we can achieve much better performance after applying our EM method and the mean error is 1.49 degree, which is shown in the last two columns.

6.2. Real dataset

We have captured the real dataset using the depth sensor of Kinect V2 and a color camera with resolution of 1920×1080 . There are several models captured, namely

the Backpack, Turtle, Book, etc.

Fig. 5 shows the comparison results of a colorful Backpack and Turtle model respectively. Fig. 5(a2) and (b2) displays the mesh model acquired from KinectFusion. As we can see the surface details are not revealed as restricted by the resolution and accuracy of the Kinect depth sensor. After applying the shading refinement [27] on the fused reference frame, some surface details are recovered, while some textures are hallucinated as geometry details as well

(Fig. 5(a3) and (b3)). Fig. 5(a4) and (b4) shows our results without applying our local match refinement step. The uneven surface in some part as marked in red is caused by misalignment. Fig. 5(a5,b5) and Fig. 5(a6,b6) displays our final results of recovered meshes and surface normals. For the Backpack, it actually experiences non-rigid deformation during the movement and therefore we only consider the front part of the backpack which is mostly rigid. We have marked some colorful patterns on the Turtle surface to make the texture more complex to show the superior performance of our pixel-wise recovery method.

To further validate our robustness against texture copy problem, the results of a colorful Book cover are demonstrated in Fig. 6. As displayed in Fig. 6(b) the textures have been successfully factored out from the image and the recovered model keeps as a planar surface after the enhancement. In comparison, the result from shading refinement method(Fig. 6(c)) is affected by the texture copy effect with fake geometric details appeared.

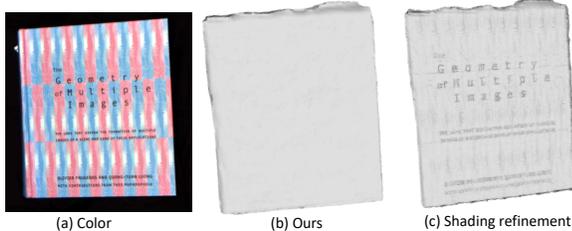


Figure 6: Results on Book Model. (a) is the reference color image. The recovered mesh surface from our method is displayed in (b). (c) shows the refined mesh with shading refinement method.

Intrinsic Image Decomposition In order to show the performance of our method in albedo recovery, we have also made some comparisons with two state-of-the-art intrinsic image decomposition approaches as displayed in Figure 7. For these two compared methods, they take the RGB-D images of the reference frame as input, as displayed in the first column. The second column shows the result from Chen [9]. The shading image is over smoothed and the geometry details are decomposed into albedo map incorrectly. The method from Jeon [17] has better results on recovered shading images for the Turtle model as displayed in the third column. However, some textures still stay at the shading image especially for the Backpack. In comparison, with our pixel-wise albedo computation method, we are able to recover a much sharper albedo map and the "texture copy" effect in the geometry is barely noticeable.

7. Conclusion

In this paper, we present a novel approach to recover surface details and its albedo map from an RGB-D video

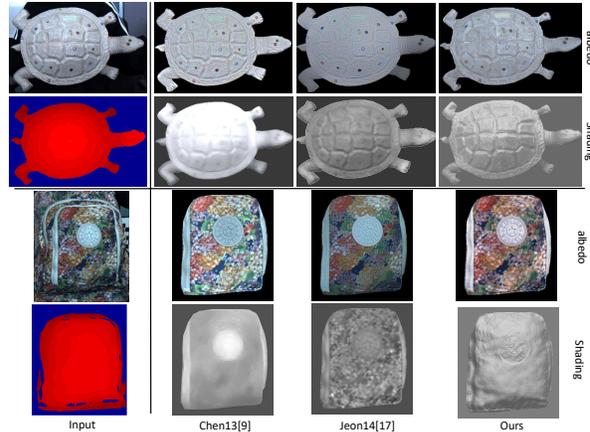


Figure 7: Comparison results on albedo recovery or intrinsic decomposition of the Turtle and Backpack model. The first column is the input color image with its depth map. The second column shows the result of Chen [9]. The third column is the decomposed albedo and shading images from method in [17]. Finally, the last column demonstrates the result achieved by our method.

sequence. The object is experiencing casual motion from which the induced illumination variation provides us the cue to recover the surface normal and its albedo as well. A robust lighting insensitive local match strategy is proposed to establish correct correspondences along the sequence. Then, the environmental lighting is estimated by exploiting the whole sequence to get rid of effect of albedos. Finally, the surface normal and its albedo is calculated robustly with our EM framework. We have validated our method on both synthetic and real dataset and compared with some state-of-the-art surface refinement and intrinsic decomposition methods. As demonstrated in the experiments, we have achieved good performance on both surface details recovery and intrinsic decomposition.

The reconstructed object is limited to objects with primarily diffuse surface. As a future work, we could extend our idea into objects that have a great portion of non-Lambertian reflection. Right now we are mainly focusing on depth enhancement, while we would like to implement all these procedure in full 3D space.

Acknowledgments

This work is partially supported by the US NFS (IIS-1231545, IIP-1543172), US Army Research grant W911NF-14-1-0437, NSFC (No. 61332017, 51475373, 61603302, 51375390), Key Industrial Innovation Chain of Shaanxi Province Industrial Area (2015KTZDGY04-01, 2016KTZDGY06-01), the fundamental Research Funds for the Central Universities (No. 3102016ZY013). Jiangbin Zheng and Ruigang Yang are the co-corresponding authors for this paper.

References

- [1] J. T. Barron and J. Malik. High-frequency shape and albedo from shading using natural image statistics. In *CVPR*, pages 2521–2528, 2011. 2
- [2] J. T. Barron and J. Malik. Shape, illumination, and reflectance from shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(8):1670–1687, 2015. 2
- [3] J. T. Barron and J. Malik. Intrinsic scene properties from a single rgb-d image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(4):690–703, 2016. 2
- [4] H. G. Barrow and J. M. Tenenbaum. Recovering intrinsic scene characteristics from images. *Comput. Vis. Syst.*, 1978. 2
- [5] S. Bi, X. Han, and Y. Yu. An l1 image transform for edge-preserving smoothing and scene-level intrinsic decomposition. *ACM Trans. Graph.*, 34(4):78, 2015. 2
- [6] M. Böhme, M. Haker, T. Martinetz, and E. Barth. Shading constraint improves accuracy of time-of-flight measurements. *Computer vision and image understanding*, 114(12):1329–1335, 2010. 1
- [7] T. Brox and J. Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE transactions on pattern analysis and machine intelligence*, 33(3):500–513, 2011. 4
- [8] A. Chatterjee and V. M. Govindu. Photometric refinement of depth maps for multi-albedo objects. In *CVPR*, pages 933–941, 2015. 2
- [9] Q. Chen and V. Koltun. A simple model for intrinsic image decomposition with depth cues. In *ICCV*, pages 241–248, 2013. 2, 8
- [10] C. H. Esteban, G. Vogiatzis, and R. Cipolla. Multiview photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3):548–554, 2008. 1
- [11] D. Ferstl, C. Reinbacher, R. Ranftl, M. Rother, and H. Bischof. Image guided depth upsampling using anisotropic total generalized variation. In *ICCV*, pages 993–1000, 2013. 1
- [12] G. Finlayson and R. Xu. Illuminant and gamma comprehensive normalisation in logrgb space. *Pattern Recognition Letters*, 24(11):1679–1690, 2003. 4
- [13] E. Garces, A. Munoz, J. Lopez-Moreno, and D. G. Diego. Intrinsic images by clustering. In *Computer Graphics Forum*, volume 31, pages 1415–1424, 2012. 2
- [14] Y. Han, J. Lee, and I. Kweon. High quality shape from a single rgb-d image under uncalibrated natural illumination. In *ICCV*, pages 1617–1624, 2013. 2
- [15] S. M. Haque, A. Chatterjee, and V. M. Govindu. High quality photometric reconstruction using a depth camera. In *CVPR*, pages 2283–2290, 2014. 2
- [16] B. K. P. Horn. *Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View*. PhD thesis, MIT, 1970. 2
- [17] J. Jeon, S. Cho, X. Tong, and S. Lee. Intrinsic image decomposition using structure-texture separation and surface normals. In *ECCV*, pages 218–233, 2014. 2, 8
- [18] K. Kim, A. Torii, and M. Okutomi. Joint estimation of depth, reflectance and illumination for depth refinement. In *ICCV Workshops*, 2015. 2
- [19] N. Kong, P. Gehler, and M. Black. Intrinsic video. In *ECCV*, pages 360–375, 2014. 2
- [20] P. Laffont and J. Bazin. Intrinsic decomposition of image sequences from local temporal variations. In *ICCV*, pages 433–441, 2015. 2
- [21] P. Laffont, A. Bousseau, S. Paris, F. Durand, and G. Dretakis. Coherent intrinsic images from photo collections. *ACM Trans. Graph.*, 31(6), 2012. 2
- [22] A. Lakdawalla and A. Hertzmann. Shape from video: Dense shape, texture, motion and lighting from monocular image streams. In *Proceedings of the First International Workshop on Photometric Analysis For Computer Vision*. INRIA, 2007. 1
- [23] E. H. Land and J. J. McCann. Lightness and retinex theory. *JOSA*, 1971. 2
- [24] K. J. Lee, Q. Zhao, X. Tong, M. Gong, S. Izadi, S. U. Lee, P. Tan, and S. Lin. Estimation of intrinsic image sequences from image+ depth video. In *ECCV*, pages 327–340, 2012. 2
- [25] D. Min, J. Lu, and M. N. Do. Depth video enhancement based on weighted mode filtering. *IEEE Transactions on Image Processing*, 21(3):1176–1190, 2012. 1
- [26] R. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *ISMAR*, pages 127–136, 2011. 3
- [27] R. Or-El, G. Rosman, A. Wetzler, R. Kimmel, and A. M. Bruckstein. Rgb-d-fusion: Real-time high precision depth recovery. In *CVPR*, pages 5407–5416, 2015. 1, 2, 6, 7
- [28] J. Park, H. Kim, Y. Tai, M. S. Brown, and I. Kweon. High-quality depth map upsampling and completion for rgb-d cameras. *IEEE Transactions on Image Processing*, 23(12):5559–5572, 2014. 1
- [29] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. In *SIGGRAPH*, pages 497–500. ACM, 2001. 2
- [30] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. *ACM transactions on graphics*, 23(3):309–314, 2004. 6
- [31] B. Shi, K. Inose, Y. Matsushita, P. Tan, S. Yeung, and K. Ikeuchi. Photometric stereo using internet images. In *3DV*, pages 361–368, 2014. 2
- [32] D. Simakov, D. Frolova, and R. Basri. Dense shape reconstruction of a moving object under arbitrary, unknown lighting. In *ICCV*, volume 3, pages 1202–1209, 2003. 1
- [33] M. F. Tappen, W. T. Freeman, and E. H. Adelson. Recovering intrinsic images from a single image. *IEEE TPAMI*, 27(9):1459–1472, 2005. 2
- [34] C. Ti, R. Yang, J. Davis, and Z. Pan. Simultaneous time-of-flight sensing and photometric stereo with a single tof sensor. In *CVPR*, pages 4334–4342, 2015. 2
- [35] C. Wu, Y. Liu, Q. Dai, and B. Wilburn. Fusing multiview and photometric stereo for 3d reconstruction under uncalibrated illumination. *IEEE transactions on visualization and computer graphics*, 17(8):1082–1095, 2011. 2

- [36] C. Wu, K. Varanasi, Y. Liu, H. Seidel, and C. Theobalt. Shading-based dynamic shape refinement from multi-view video under general illumination. In *ICCV*, pages 1108–1115, 2011. [2](#)
- [37] C. Wu, M. Zollhöfer, M. Niener, M. Stamminger, S. Izadi, and C. Theobalt. Real-time shading-based refinement for consumer depth cameras. *ACM Transactions on Graphics (TOG)*, 33(3), 2014. [1](#), [2](#)
- [38] T. Wu and C. Tang. Photometric stereo via expectation maximization. *IEEE transactions on pattern analysis and machine intelligence*, 32(3):546–560, 2010. [2](#), [5](#)
- [39] T. Wu, K. Tang, C. Tang, and T. Wong. Dense photometric stereo: A markov random field approach. *IEEE transactions on pattern analysis and machine intelligence*, 28(11):1830–1846, 2006. [2](#)
- [40] J. Yang, X. Ye, K. Li, C. Hou, and Y. Wang. Color-guided depth recovery from rgb-d data using an adaptive autoregressive model. *IEEE Transactions on Image Processing*, 23(8):3443–3458, 2014. [1](#)
- [41] L. Yu, S. Yeung, Y. Tai, and S. Lin. Shading-based shape refinement of rgb-d images. In *CVPR*, pages 1415–1422, 2013. [1](#), [2](#)
- [42] L. Zhang, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and motion under varying illumination: Unifying structure from motion, photometric stereo, and multiview stereo. In *ICCV*, pages 618–625, 2003. [1](#)
- [43] Q. Zhang, M. Ye, R. Yang, Y. Matsushita, B. Wilburn, and H. Yu. Edge-preserving photometric stereo via depth fusion. In *CVPR*, pages 2472–2479, 2012. [2](#), [6](#)
- [44] Q. Zhao, P. Tan, Q. Dai, L. Shen, E. Wu, and S. Lin. A closed-form solution to retinex with nonlocal texture constraints. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(7):1437–1444, 2012. [2](#)
- [45] Q. Zhou and V. Koltun. Color map optimization for 3d reconstruction with consumer depth cameras. *ACM Transactions on Graphics (TOG)*, 33(4):155, 2014. [5](#)
- [46] Z. Zhou, Z. Wu, and P. Tan. Multi-view photometric stereo with spatially varying isotropic materials. In *ICCV*, pages 1482–1489, 2013. [2](#)