

Shape and Material Capture at Home

Daniel Lichy¹ Jiaye Wu¹ Soumyadip Sengupta² David W. Jacobs¹

¹University of Maryland, College Park ²University of Washington

dlichy@umd.edu, jiayewu@umiacs.umd.edu, soumya91@cs.washington.edu, djacobs@cs.umd.edu

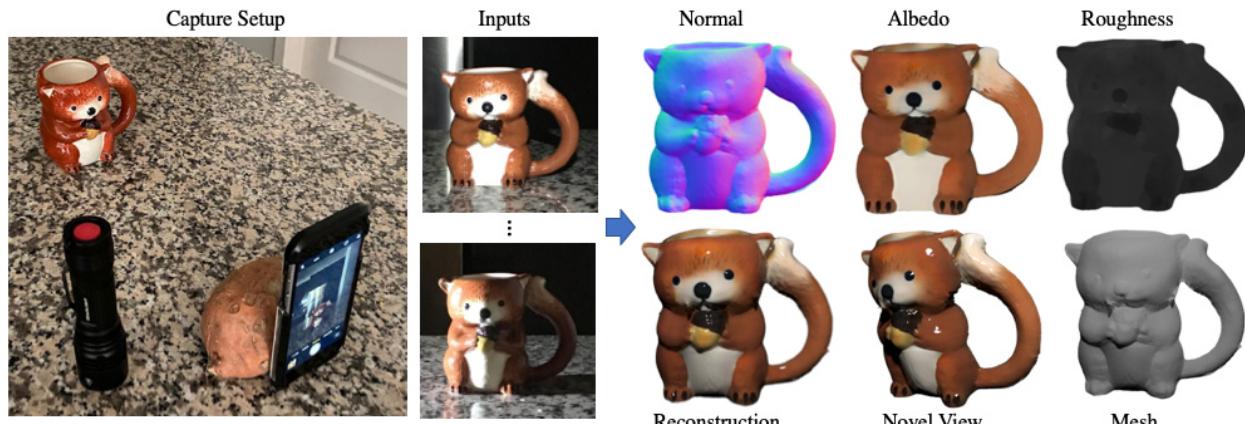


Figure 1: A bare-bones capture setup (left) consisting of only an iPhone, a flashlight and a sweet potato. The user captures a few images (with voice control) of the object by moving around with a flashlight directed at the object. We then reconstruct surface normal, albedo and roughness at high-resolution (1024×1024). Normals are integrated to create a mesh that is rendered with novel views and illuminations.

Abstract

In this paper, we present a technique for estimating the geometry and reflectance of objects using only a camera, flashlight, and optionally a tripod. We propose a simple data capture technique in which the user goes around the object, illuminating it with a flashlight and capturing only a few images. Our main technical contribution is the introduction of a recursive neural architecture, which can predict geometry and reflectance at $2^k \times 2^k$ resolution given an input image at $2^k \times 2^k$ and estimated geometry and reflectance from the previous step at $2^{k-1} \times 2^{k-1}$. This recursive architecture, termed RecNet, is trained with 256×256 resolution but can easily operate on 1024×1024 images during inference. We show that our method produces more accurate surface normal and albedo, especially in regions of specular highlights and cast shadows, compared to previous approaches, given three or fewer input images.

1. Introduction

Capturing an object’s shape and material is a long-studied problem in Computer Vision and Graphics, with a broad range of applications such as Augmented and Virtual Reality. With the rise of e-commerce, it is now even more important to develop a system that allows sellers to capture 3D shape and material of their product with relative ease. It

can also be extremely useful for digital artists who can use captured 3D objects as a starting point for their models.

In this work, we present a system where the shape and reflectance of an object can be captured by a user with everyday household items. We propose a capture setup that requires a camera, flashlight, and ideally a tripod. The user can capture one or multiple images of an object by illuminating the object with a flashlight from multiple directions. This relatively straight-forward data capture setup provides us with important photometric cues necessary for high-quality reconstruction.

High quality object geometry can be recovered using a large number of images either from different views (Multi-View Stereo) [48, 67, 68], or from different lighting variations (Photometric Stereo) [10, 13, 15]. Our work falls in the category of both Photometric Stereo (PS) and Shape-from-Shading (SfS). PS techniques generally require a light-stage setup in a dark room to capture an object. Although these techniques often produce high-quality reconstruction given a large enough number of images, it is extremely difficult to create such a capture setup at home. On the other hand, single image based methods have shown success in capturing geometry and material using a camera with co-directional flash [37], often with an extra image without the flash [7]. However, these approaches rely on

priors learned from synthetic data with limited generalization to real data, especially with slight variation in capture setup. In this work, we aim to meet reconstruction quality and ease of capture in the middle. We propose to capture only a few images (at most six) of the object from approximately known lighting directions in a weakly calibrated fashion to achieve satisfactory reconstruction quality.

We introduce a novel, recursive architecture, RecNet, that is capable of predicting normal, albedo and roughness at high resolution. This allows our approach to create a high-quality reconstruction with fine details. RecNet is trained to predict normal, albedo and roughness at $2^k \times 2^k$ resolution given an input image at $2^k \times 2^k$ and predicted normal, albedo and roughness from the a previous iteration at $2^{k-1} \times 2^{k-1}$. We parameterize shape with surface normal and reflectance with albedo and roughness following the Cook-Torrance model. RecNet is trained with resolutions between 64×64 and 256×256 , but can then be applied to images of arbitrary resolution. To start this recursive process off, we use a small network, InitNet, that takes in the image at 32×32 and estimates the normals, albedo, and roughness at the same resolution. We demonstrate that RecNet produces high-quality results for 1024×1024 images, yielding a mesh with enough vertices to capture fine detail.

An alternative approach is to train at the same resolution as inference. However, training on higher resolution input offers its own unique challenges; it requires high-quality synthetic data and a very deep convolutional network. This would necessitate large amounts of GPU memory and training time, making it extremely difficult if not impossible. Thus PS techniques that require a large number of images (50-100) often train on pixels or patches [62, 27, 41, 10, 13, 15], but they do not perform well when using a small number of images. Most previous single image techniques can only be run at 256×256 [37, 7]. This is because photometric cues from a single image are not sufficient to determine normals using local information alone, and the global priors the network learns from 256×256 data do not transfer well to higher resolution. A network trained on lower resolution, when applied to higher resolutions effectively shrinks its receptive field, implicitly assuming that the normals are conditionally independent of the rest of the image, preventing it from capturing global context. In contrast, we assume that normals at higher resolution depend on the normals at lower resolution within the receptive field via our recursive architecture, RecNet. Normals at lower resolution reflect larger global context with the same receptive field and thus help our network to reconstruct better shape and reflectance utilizing the global context.

Our network can also handle anywhere from one to six images. When multiple images are used, we compare with a current state-of-the-art PS technique, SDPS-Net [10], and

when a single image is used as input, we compare with [37, 7], which assume flash-light co-directional with the camera. We predict higher quality surface normals compared to SDPS-Net, especially on objects with spatially varying BRDF and when given a low number of input images. While SDPS-Net only predicts geometry, we also predict material reflectance as albedo and roughness. In comparison to single image techniques, we predict better surface normals and albedo, especially in the regions of cast shadows and specular reflectance.

In summary, our contributions are as follows:

- We introduce a weakly calibrated photometric stereo technique, where shape and material of an object can be captured easily at home with a camera, flashlight, and optionally a tripod.
- We present a recursive multi-resolution architecture that can handle a varying number of input images (1-6) at arbitrary resolution, which can be higher than the training resolution (training at 256, inference at 1024) with minimal artifacts.

2. Prior Work

Photometric Stereo (PS), first introduced in [65], aims to reconstruct the shape of an object from images under varying illumination, and is a long studied problem in computer vision (see [2, 19, 60] for surveys). When intensity and direction of the lights are unknown the problem is referred to as Uncalibrated PS (UPS) [4], which is often more challenging due to ambiguity [5]. Previous approaches either assume Lambertian reflectance [21, 51, 25] or aim to model non-diffuse reflectance using various BRDF models [45, 22, 50, 44]. A few works [26, 49, 8, 30] aim to solve PS for more general environment lighting.

With the recent success of deep learning, researchers have shown renewed interest in pushing the boundaries of PS. Recent works [56, 27, 41] focus on the calibrated setting with a large number of input images and learn on a per-pixel basis, thus ignoring global context. Taniai *et al.* [62] proposed a network that is trained per object to predict normal and BRDF using unsupervised reconstruction loss with known lighting. Chen *et al.* [14] trained a network on synthetic data to predict normals with known lighting. For UPS, recent works focus on predicting the lighting and then using it to predict normals [11], often in a recursive fashion [15], and can handle spatially varying BRDF [12]. While most PS methods require a large number of images (50-100), researchers have also attempted to solve PS with few images [59, 46, 52, 35, 71]. Researchers have also proposed techniques to solve PS with known direction and unknown intensity [17, 53, 52], and with inaccurate direction and intensity [54]. In this work, our system assumes that the lighting direction is known to within 10° to 15° , and the intensity is unknown. It requires only a few images (1-6) and predicts both normal and spatially-varying BRDF.



Figure 2: Our at-home capture setup uses a camera, a flashlight, and optionally a tripod and a remote. A user moves around the object to capture a single or multiple images with variation in lighting. Images are captured from the ‘Front’ and the five positions shown above.

While PS usually requires a large number of observations, another research direction, Shape from Shading [70], attempts to predict shape and reflectance from a single image. Previous approaches like Barron and Malik [3] relied on extensive, manually designed shape priors. With recent advances in deep networks such priors are often learned from large volumes of synthetic data. Recent works made progress in solving Shape from Shading for faces [32, 58, 34] and Inverse Rendering for general scenes [69, 57, 36], all from a single image. For a generic object, researchers have attempted to estimate reflectance and illumination from a single image [42, 43, 23, 24, 55, 47]. Liu *et al.* [40] predicted normal and DS-BRDF from a single image by training on ShapeNet objects. Li *et al.* [37] proposed a framework where a single image of an object is captured with a flash co-directional with the camera and predicted normal, depth, spatially varying BRDF and illumination. They trained a cascaded architecture on synthetic data, generated similarly to our approach. An extension of single image based shape from shading is to capture two images, with and without a co-directional flash, as presented in Boss *et al.* [7], which also uses a cascaded architecture. Both [37, 7] can only make predictions at 256×256 resolution. In contrast, we present a unified architecture that predicts state-of-the-art normal and spatially varying BRDFs from a single image or from a few images (2-6) illuminated with a flashlight and ambient light (in at least one image the flashlight is approximately co-directional with the camera), at 1024×1024 resolution.

A key contribution of our technique is a recursive multi-scale network, where a single network is trained with multiple resolutions. The importance of multi-scale information in normal estimation from a single image goes back well before the age of deep learning. This is due to ambiguities that cannot be resolved with local information alone but require global context. Separate networks for local and global context and a fusion step to exploit multi-scale information were used in [63, 20]. Recent deep networks often aim to achieve this by using a cascaded architecture where the later stages in the cascade are responsible for learning details [37, 36], but all stages operate at the same resolution. However, these methods trained for 256×256 cannot operate on 1024×1024 , as they involve fully connected layers. Often optimization is also used for refining the network prediction [6] at the cost of runtime. Researchers have also explored cascaded architectures that operate at different resolutions

in segmentation [38, 39], super-resolution [33], and generative models [16, 31]. In contrast to these methods, our network uses the same weights at each step for predicting normal, albedo and roughness at resolution $2^k \times 2^k$ given an input image at $2^k \times 2^k$ and predicted normal, albedo and roughness at $2^{k-1} \times 2^{k-1}$. This also enables us to operate on images larger than those we train on.

3. Our Approach

Our goal is to capture high-quality shape and reflectance of household objects with minimal effort. To achieve this goal, we make several design choices. First, we parameterize shape with surface normals, which will then be integrated into a depth map and converted into a mesh representation. Second, we parameterize spatially-varying BRDF (SV-BRDF) with albedo and roughness using the Cook-Torrance BRDF model [18] (for details please see the supplement). We aim to reconstruct surface normal, albedo and roughness at 1024×1024 resolution, in contrast to 256×256 used in [7, 37], such that the reconstructed mesh can preserve details. Finally, our capture setup includes a camera mounted on an inexpensive tripod about 2' from the object. Then the user captures multiple images of the object by walking around with a flashlight, with at least one image where the light is roughly co-directional with the camera. Our data capture setup inspires us to generate synthetic data reflective of our real data.

3.1. Our Capture Setup

We want our capture setup to be simple such that, with instructions, a user can easily replicate it in their home, and also provide enough distinct observations for high-quality reconstruction. For our capture setup we used the following items (i) A CANON 2000D DSLR camera; (ii) a tripod; (iii) a remote trigger; and (iv) an LED flashlight. As we show in Figure 1 our capture setup can be performed with just a camera phone, stand, and flashlight.

Our capture protocol is as follows. The object is placed on a flat surface in a room dimly lit relative to the flashlight (*dark room not required*). The camera is pointed toward the object; its optical axis making approximately a 20 to 40° angle with the surface. The user can then take one to six images by moving around with a flashlight pointed at the object. In one image, the light should be approximately co-directional with the camera. For six images, the

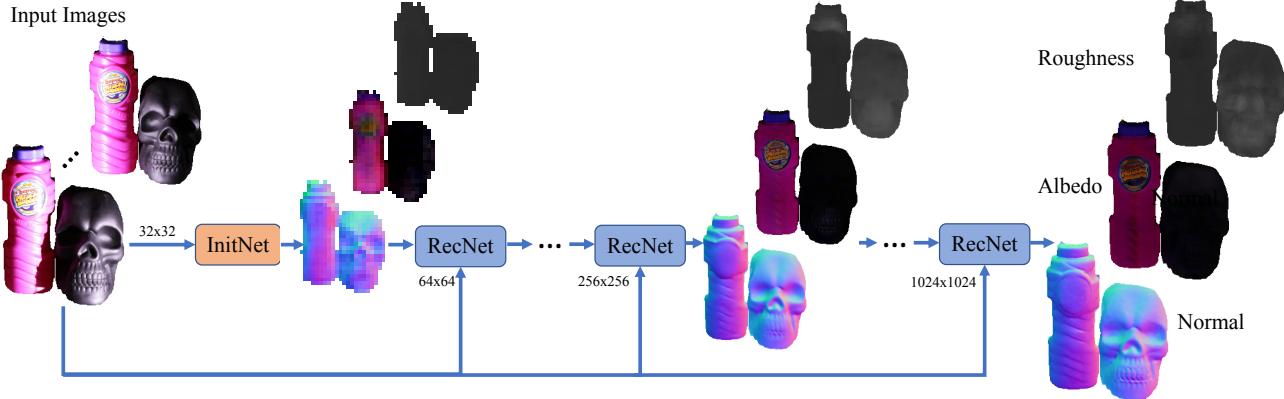


Figure 3: We propose a recursive multi-resolution architecture, RecNet, that predicts surface normal, albedo and roughness from the input image(s) and from the prediction at the previous step by continuously upsampling by a factor of 2. The recursion is initialized by InitNet.

flashlight is pointed at the object from the following directions: right ($+90^\circ$), front-right ($+45^\circ$), co-directional with the camera (0°), front-left (-45°), left (-90°), and directly above. Note that these directions are weakly calibrated, i.e. our system roughly knows the location of the light and can tolerate $\pm 15^\circ$ error in angular direction. We chose these directions because we felt they are easy for the user to estimate. Please see Fig 2, which depicts a user capturing an object following our capture setup. To evaluate our method, we capture a large real dataset of 111 scenes containing a variety of objects we obtained at various discount stores.

3.2. Synthetic Data

Since it is extremely difficult to capture real-world objects with ground-truth shape and reflectance, we rely on synthetic data to train our network. Our synthetic data generation is motivated by our real-world data capture setup. We perform synthetic data generation with primitive shapes, similar to Xu *et. al* [66]. We consider 5,000 objects, each consisting of 1-9 primitive shapes (cubes, ellipsoids, and cylinders) with surface perturbations of various frequencies. Additionally, we consider 14 realistic shapes from the Sculptures dataset [64].

For spatially varying albedo we use free textures from [1]. Unlike [66, 37, 7] we do not use the roughness supplied with the textures because their details are too small to be recovered faithfully. Instead, we apply a random roughness to each primitive shape in the scene. For sampling roughness, we first sample a Phong exponent from an exponential distribution that is approximately the same as the one used by [47]. We then convert it into a Beckmann equivalent roughness [29]. We observe that this design choice forces the network trained on this data to rely more on photometric cues than prior associations between albedo and roughness. Thus the network can generalize better to real data. More details on our synthetic data generation procedure can be found in the supplementary material.

3.3. Recursive Multi-Resolution Network

Our goal is to predict high-resolution surface normals, albedo and roughness, such that when we integrate the normals to create a mesh, we can preserve the details of the object. An obvious choice is to train a convolutional network, e.g. ResNets [72], on high-resolution input-output pairs. However, this is extremely difficult to achieve for the following reasons. First, the primitive shapes used for data generation lack details, and so do most synthetic shape datasets like Shapenet [9]. Thus the synthetic high-resolution data will lack the details often observed in real-world objects. Second, training a ResNet with sufficient depth on 1024×1024 images will require significant memory and training time, which often makes it impossible to train on a single GPU with batch-size of one. ResNets trained on low-resolution 256×256 images introduce additional artifacts when tested on 1024×1024 images. Thus our goal is to design a network that can be trained at low-resolution yet still produce satisfactory results on high-resolution data during inference.

The input to our system is a set of images of an object illuminated by a light co-directional with the camera and optionally up to five additional images of the object illuminated by directional light from the right, center-right, center-left, left, and above. We also input a segmentation mask. Then our system predicts surface normal N , albedo A and roughness R . We use an interactive segmentation tool [61] to create a the mask, usually with just 3-5 clicks.

Our network, depicted in Figure 3, consists of two components, an initialization network (InitNet) and a recursive network (RecNet). InitNet receives input at a down-sampled resolution of 32×32 , I^{32} , and predicts normal, albedo, and roughness at the same resolution. This is used as initialization to the recursive network RecNet, which continuously upsamples the result by a factor of 2. RecNet starts by taking in the 32×32 normal, albedo, and roughness predictions made by InitNet along with the input images

downsampled to 64×64 and predicts the normal, albedo, and roughness at 64×64 . These are then recursively fed into RecNet along with the inputs downsampled to 128×128 . This procedure is repeated until the normal, albedo, and roughness match the resolution of the original input images. This procedure is described in pseudo-code in Algorithm 1.

Algorithm 1 Network procedure: The input image at $2^K \times 2^K$ is down sampled to $2^k \times 2^k$ resolution, denoted by I^{2^k} , for k from K to 6. $A^{2^k}, N^{2^k}, R^{2^k}$ represent albedo, normal and roughness predicted at $2^k \times 2^k$ resolution, respectively.

```

 $N^{2^5}, A^{2^5}, R^{2^5} = \text{InitNet}(I^{2^5})$ 
for  $k = 6$  to  $K$  do
     $N^{2^k}, A^{2^k}, R^{2^k} = \text{RecNet}(I^{2^k}, N^{2^{k-1}}, A^{2^{k-1}}, R^{2^{k-1}})$ 
end for

```

InitNet consists of three small Resnets, one for albedo, normal, and roughness, that only differ in the number of output channels: 3 for albedo; 2 for normal (we only predict the x and y component of the normal, z is calculated as $\sqrt{1 - x^2 - y^2}$); and 1 for roughness. The ResNets in InitNet have no downsampling and only two residual blocks. The input to the network is a concatenation of six or fewer images and the mask of the object. When less than six images are used, respective channels in the input are zeroed. The RecNet architecture also consists of 3 ResNets; each ResNet contains 8 residual blocks, at its narrowest point each feature resolution is 1/4 the input resolution. The input to RecNet is six images and the mask along with the predicted normal, albedo and roughness from the previous step upsampled by a factor of 2. Please see the supplement for a detailed description of the network architectures.

Our architecture at test time applies the same network at each step to refine the network predictions from the previous step by a scale factor of 2. Thus we aim to train this network to be scale-independent. All networks are trained simultaneously. The networks are applied as in algorithm 1. This produces estimates of albedo, normal, and roughness at four resolutions from 32 to 256. We take the L_1 loss of normal, albedo, and roughness at all four resolutions and sum them to get our full loss function:

$$\min_{N, A, R} \sum_{k=5}^8 L_1(A^{2^k}, \hat{A}^{2^k}) + L_1(N^{2^k}, \hat{N}^{2^k}) + L_1(R^{2^k}, \hat{R}^{2^k}),$$

where \hat{N} , \hat{A} and \hat{R} denotes ground-truth normal, albedo, and roughness. Thus, RecNet is being trained at 3 scales simultaneously, and InitNet is only being trained at 32×32 .

To explain the necessity of this recursive architecture, consider the rendering of a pyramid in Figure 4A. The pyramid is diffuse with uniform albedo, and the light is co-directional with the camera. If we would only consider a local window around point p we would just see a uniform color; it requires global context to disambiguate the normal

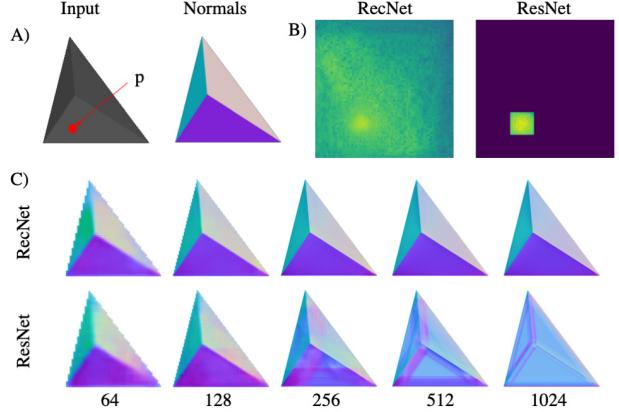


Figure 4: (A) A pyramid rendered with co-directional light and camera, and its ground truth normals. (B) Visualization of the gradient of the input image with respect to output pixel p , for RecNet and ResNet at 512×512 resolution. Both RecNet and ResNet are trained at 256×256 and have receptive fields 93 and 85, respectively. Note that for the ResNet the output at pixel p only depends on the input in a neighborhood of pixel p , this is not the case for RecNet. (C) Normal prediction by RecNet and ResNet at different input resolutions.

at this location. A fully convolutional architecture, such as a Resnet, has a finite receptive field; thus if we have this image at very high resolution, the network's receptive field will not be large enough to capture the global context necessary to disambiguate the normal. Furthermore, even if we trained a network with a very large receptive field, it seems unlikely that it could learn relationships more distant than the size of the training data itself, in our case 256×256 .

In Figure 4C we see that as the resolution of the triangle grows, it requires more and more relatively distant information to predict the normal correctly; thus the quality of the prediction by the ResNet deteriorates. In contrast, with the RecNet the receptive field is proportional to the number of iterations i.e. image size; thus its prediction does not deteriorate with increasing image resolution. We show this in 4B by plotting the gradient of the input image with respect to pixel p . We observe that for the ResNet the normal at pixel p is determined only by a small neighborhood around p whereas the RecNet also takes into account pixels a much greater distance away. Although this example seems contrived, we show in section 5 that for real data, the ResNet produces significantly more artifacts than the RecNet.

3.4. Implementation

We implement our network in Pytorch. We train with the Adam optimizer for 50 epochs with a batch size of 10 and a learning rate of 10^{-4} . In each batch, one to five images out of six are randomly zeroed such that the probability of getting 1,2,3,4,5,6 non-zero images is 0.3, 0.2, 0.2, 0.1, 0.1, 0.1. Note the co-directional light image is never ze-

roed. Training took approximately two days on two Nvidia RTX2080Ti GPUs.

After our network predicts surface normal, albedo and roughness, we integrate the normals to obtain a depth map. Then we consider each pixel of the depth map as a vertex of a mesh. For details on our integration procedure, please see the supplement.

4. Experiments

We compare with various state-of-the-art approaches for shape and material estimation:

- SPDS-Net [10] is a PS technique that only predicts surface normals from multiple images. Although not emphasized in the original paper, it can also be used for predicting normals from a single image. SDPS-Net, like ours, can predict normals at 1024×1024 resolution.
- Li'18 [37] predicts surface normal, albedo and roughness from a single image captured with co-directional camera and light at 256×256 resolution.
- Boss'20 [7] also predicts normal, albedo and roughness but requires two images, with and without flash, at 256×256 resolution.

Since it is extremely difficult to capture ground-truth albedo and roughness of household objects, our evaluation relies on qualitative comparison. However, it is possible to quantitatively evaluate normals on real data with the DiLiGenT dataset [60]. DiLiGenT consists of 10 objects with diverse materials captured under 96 calibrated lighting directions with ground truth normals. The lighting directions lie approximately on a rectangle and are all within roughly 45 degrees of the camera.

Surface Normal estimation with multiple images. In Table 1, we compare our normal estimation with SDPS-Net on the DiLiGenT dataset with three images coming from the front, front-right, and front-left. We are restricted to at most three images on this dataset because images from the left, right, and above are not included. We report a mean angular error (MAE) of 12.5° compared to 23.3° for [10]. However, SDPS-Net is trained on 32 images, so we retrain it with only three images as input. Although this improves the quality of SDPS-Net our method is still numerically superior to it. Furthermore, the qualitative evaluation in Figure 5, clearly shows the superior performance of our method on challenging objects. We also show superior performance to SDPS-Net on the two image task (see supplementary).

Surface Normal estimation with single image. In Table 2 we compare normal estimation error of our method with that of Li'18 [37] and SDPS-Net [10] on DiLiGenT using a single image captured with approximately co-directional light and camera. We also retrained SDPS-Net just on a single image, although our method is trained to work with up to 6 images at once, and still, we outperform SDPS-Net. We could not compare with Boss'20 [7] on DiLiGenT as it requires two images with and without flash. In Figure 6, we present a qualitative comparison of our

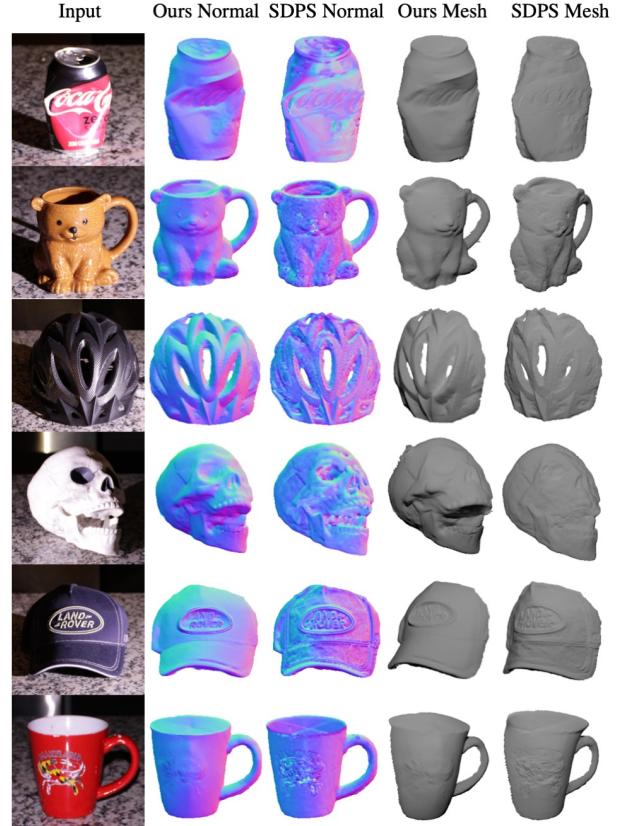


Figure 5: Comparison of our normal estimation method with that of SDPS-Net retrained on 3 images. Tested with front, front-right and front-left images.

method with that of SDPS-Net retrained on a single image, Li'18 and Boss'20 on images captured by us. Both quantitative and qualitative evaluations show that our method outperforms state-of-the-art normal estimation techniques for single and few image inputs.

Material Estimation. We present a visual comparison of albedo in Figure 7 and roughness in Figure 8 estimated by our method with that of Li'18 using a single image and Boss'20 with the same image and an additional image without directional light. We also show the albedo and roughness predicted by our method with six images as input. We observe that our method is often better in removing cast shadows and specular highlights from the albedo. Roughness is challenging for all methods. We often observe that even relative roughness predictions do not agree between methods. More qualitative results on our data and data collected by [37] can be found in the supplement.

5. Ablation Studies

Evaluation of Recursive Architecture. To evaluate the effectiveness of our recursive architecture, we train a ResNet based on [72, 28]. Our recursive network has approximately 14.5M parameters, and the vanilla Resnet has

Table 1: **Three Image Results on DiLiGenT** Comparison of SDPS-Net [10], SDPS-Net retrained with 3 input images, and our method on DiLiGenT. Lights are from the front, front-right, and front-left of the camera. MAE (in degrees) for each object is reported.

Algorithm	ball	cat	pot1	bear	pot2	buddha	goblet	reading	cow	harvest	mean
SDPS-Net [10]	15.4	22.4	25.7	17.9	18.5	24.0	31.0	29.3	21.4	27.1	23.3
SDPS-Net (3 image retrained)	5.7	15.3	13.1	7.6	13.0	18.7	24.6	22.5	12.1	22.3	15.5
ResNet (ablation)	5.8	14.7	12.9	8.3	12.3	14.9	23.4	17.3	15.7	23.2	14.8
Ours	5.7	12.7	10.6	7.1	10.2	13.9	17.0	16.9	9.9	20.6	12.5

Table 2: **Single Image Results on DiLiGenT** Comparison of [10], SDPS-Net retrained with 1 image (fully calibrated), [37] and our method on DiLiGenT using a single input image with light approximately co-directional with the camera. MAE (in degrees) is reported.

Algorithm	ball	cat	pot1	bear	pot2	buddha	goblet	reading	cow	harvest	mean
SDPS-Net [10]	36.0	35.4	36.3	34.2	36.7	44.3	43.2	43.4	35.4	42.0	38.7
SDPS-Net (1 image retrained)	6.0	21.0	17.8	11.2	18.1	26.7	27.8	29.0	17.3	34.1	20.9
Li'18 [37]	20.4	29.7	19.5	27.2	20.2	32.1	22.6	32.3	21.4	37.1	26.3
ResNet (ablation)	5.4	22.8	16.4	9.2	15.1	23.6	27.8	24.6	15.5	30.5	19.1
Ours	7.1	21.1	12.7	8.3	12.7	20.7	20.3	22.3	11.7	29.9	16.7

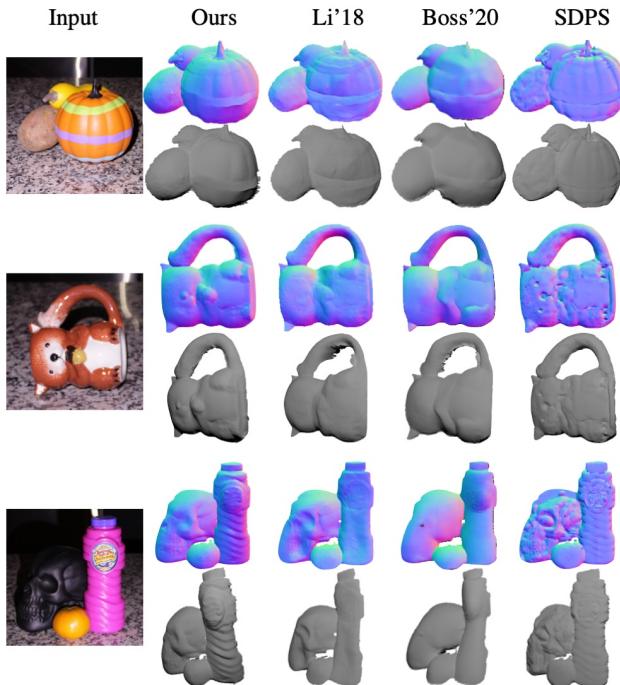


Figure 6: Comparison of normal estimation with Li'18 [37], Boss'20 [7] and SDPS-Net retrained on a single image [10].

approximately 16M parameters. Both our recursive network, RecNet, and ResNet are trained at 256×256 resolution and tested on 1024×1024 . Tables 1 and 2 show quantitatively that our recursive architecture outperforms the ResNet on DiLiGenT in the 3 image and 1 image cases, respectively. Figure 9 demonstrates ResNet introducing artifacts when tested at 1024×1024 .

Training with a Fixed Number of Inputs. We also trained specific variants of our network that only use a fixed number of images as input, either one or three, rather than a random number of inputs during training. This means that instead of having one network that can handle any number



Figure 7: Comparison of albedo estimation between Li'18 [37], Boss'20 [7], our method with one image (ours single) and our method with six images (ours six).

of inputs, we obtain separate network weights for different numbers of input images. We found that our proposed training method is only slightly worse ($< 2^\circ$) on DiLiGenT than networks trained separately for 1 and 3 image inputs. This is shown in Table 3. However, we felt that having a single network that performs well across different numbers of input images is more desirable than having different networks for different numbers of input images.

Robustness to Angle and Intensity Variation. Although we need lighting to be weakly calibrated, our method is quite robust to variation in capture directions. We

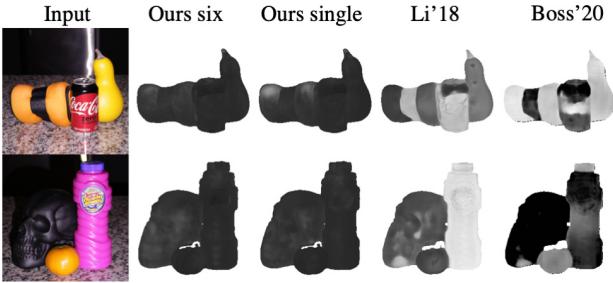


Figure 8: Comparison of roughness estimation between Li’18 [37], Boss’20 [7], our method with one image and our method with six images. Brighter indicates rougher (less specular).

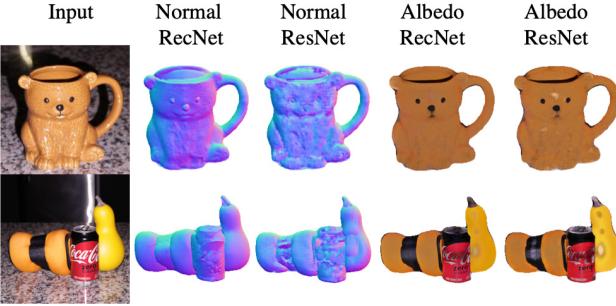


Figure 9: Illustration of artifacts produced by ResNet when trained at 256×256 and tested at 1024×1024 compared to RecNet.

	3 image input		1 image input	
	Ours	Ours-3	Ours	Ours-1
MAE (in degrees)	12.5	10.7	16.7	16.1

Table 3: Our network can handle any number of input images (1-6), and performs comparable to training separate networks with a fixed number of input images; Ours-3 for 3 image input network and Ours-1 for single image input network.

test this by computing the MAE on DiLiGenT, as the input lighting directions move further away from our prescribed directions. Since our front-right and front-left directions are at the extremes of the range present in DiLiGenT, we start from these extremes and move progressively inward. We leave the center image fixed for this experiment. Table 4 shows a minimal increase in error while the lights are within about 12.5° of the optimal.

We also demonstrate our method’s robustness to intensity variation using DiLiGenT. To simulate variation in intensity, we multiply each HDR image in DiLiGenT by a random scalar sampled from a Gaussian with unit mean and various standard deviations before tonemapping and clamping between 0 and 1. We repeat the experiment five times at each standard deviation. Table 5 shows the error at each standard deviation averaged over the five runs. We see that only at a high standard deviation, above about 0.8, when many pixel values are equal to 0 and 1, do we get a significant increase in normal prediction error.

Effect of Number of Input Images. Figure 10 shows

Deviation (degrees)	4.8	8.2	12.2	16.7	21.6
MAE	12.5	12.4	12.7	13.4	15.0

Table 4: Error (MAE in degrees) on the DiLiGenT dataset as the direction of the light deviates from the optimal.

s.d. scaling	0	0.1	0.2	0.4	0.8	1.2
MAE	12.5	12.5	12.5	12.9	14.6	16.2

Table 5: Error (MAE in degrees) on the DiLiGenT dataset as the intensity changes based on standard deviation (s.d.) of a unit mean gaussian.

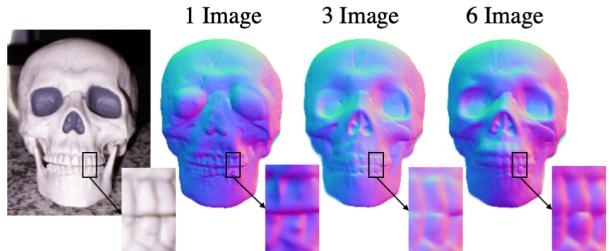


Figure 10: Results of our network with 1, 3 and 6 input images. how normal quality improves with more input images.

6. Conclusion

Shape and reflectance capture is a fundamental research problem in Computer Vision and Graphics, with applications in AR/VR, e-commerce etc. Our goal is to create a technique that allows users to capture high-quality shape and reflectance of an object in a household setting, with a camera, a flashlight, and ideally a tripod. Our technique relies on weakly calibrated capture done with a handheld flashlight. Our method is robust up to $10-15^\circ$ error in lighting direction. Weakly calibrated photometric stereo, even with error in lighting direction, appears to produce better reconstruction than uncalibrated Photometric Stereo.

Our main technical novelty is the recursive neural network RecNet, which can predict geometry and reflectance at $2^k \times 2^k$ resolution given an input image at $2^k \times 2^k$ and estimated geometry and reflectance from the previous step at $2^{k-1} \times 2^{k-1}$. This allows us to train at 256×256 and test at higher resolution, such as 1024×1024 . Training a neural network directly on higher resolution data is extremely difficult due to computational bottlenecks such as memory and training time and lack of large scale high-resolution synthetic data. Previous architectures suffer poor generalization when they train at low-resolution and perform inference at higher resolution due to a limited receptive field that does not capture global context. Our recursive architecture has a unique property in which the receptive field is doubled at each step, thus effectively allowing it to capture global information even with high-resolution images.

Acknowledgment This research is supported by the National Science Foundation under grant no. IIS-1526234 and IIS-1910132.

References

- [1] 3d textures. <https://3dtextures.me/>. Accessed: 2020.
- [2] Jens Ackermann and Michael Goesele. A survey of photometric stereo techniques. *Foundations and Trends® in Computer Graphics and Vision*, 9(3-4):149–254, 2015.
- [3] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence*, 37(8):1670–1687, 2014.
- [4] Ronen Basri, David Jacobs, and Ira Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of computer vision*, 72(3):239–257, 2007.
- [5] Peter N Belhumeur, David J Kriegman, and Alan L Yuille. The bas-relief ambiguity. *International journal of computer vision*, 35(1):33–44, 1999.
- [6] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, David Kriegman, and Ravi Ramamoorthi. Deep 3d capture: Geometry and reflectance from sparse multi-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5960–5969, 2020.
- [7] Mark Boss, Varun Jampani, Kihwan Kim, Hendrik P.A. Lensch, and Jan Kautz. Two-shot spatially-varying brdf and shape estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [8] Mohammed Brahimi, Yvain Quéau, Bjoern Haefner, and Daniel Cremers. On the well-posedness of uncalibrated photometric stereo under general lighting. In *Advances in Photometric 3D-Reconstruction*, pages 147–176. Springer, 2020.
- [9] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [10] Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, and Kwan-Yee K. Wong. Sdps-net: Self-calibrating deep photometric stereo networks. In *CVPR*, 2019.
- [11] Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, and Kwan-Yee K Wong. Self-calibrating deep photometric stereo networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8739–8747, 2019.
- [12] Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, and Kwan-Yee Kenneth Wong. Deep photometric stereo for non-lambertian surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [13] Guanying Chen, Kai Han, and Kwan-Yee K. Wong. Ps-fcn: A flexible learning framework for photometric stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [14] Guanying Chen, Kai Han, and Kwan-Yee K Wong. Ps-fcn: A flexible learning framework for photometric stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [15] Guanying Chen, Michael Waechter, Boxin Shi, Kwan-Yee K Wong, and Yasuyuki Matsushita. What is learned in deep uncalibrated photometric stereo? In *European Conference on Computer Vision*, 2020.
- [16] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1511–1520, 2017.
- [17] Donghyeon Cho, Yasuyuki Matsushita, Yu-Wing Tai, and In So Kweon. Semi-calibrated photometric stereo. *IEEE transactions on pattern analysis and machine intelligence*, 42(1):232–245, 2018.
- [18] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Trans. Graph.*, 1(1):7–24, Jan. 1982.
- [19] Jean-Denis Durou, Maurizio Falcone, Yvain Quéau, and Silvia Tozza. Advances in photometric 3d-reconstruction, 2020.
- [20] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS’14*, page 2366–2374, Cambridge, MA, USA, 2014. MIT Press.
- [21] Paolo Favaro and Thoma Papadimitri. A closed-form solution to uncalibrated photometric stereo via diffuse maxima. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 821–828. IEEE, 2012.
- [22] Athinodoros S Georgiades. Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. In *null*, page 816. Ieee, 2003.
- [23] Stamatios Georgoulis, Konstantinos Rematas, Tobias Ritschel, Mario Fritz, Luc Van Gool, and Tinne Tuytelaars. Delight-net: Decomposing reflectance maps into specular materials and natural illumination. *arXiv preprint arXiv:1603.08240*, 2016.
- [24] Stamatios Georgoulis, Konstantinos Rematas, Tobias Ritschel, Efstratios Gavves, Mario Fritz, Luc Van Gool, and Tinne Tuytelaars. Reflectance and natural illumination from single-material specular objects using deep learning. *IEEE transactions on pattern analysis and machine intelligence*, 40(8):1932–1947, 2017.
- [25] Bjoern Haefner, Yvain Quéau, and Daniel Cremers. Photometric segmentation: Simultaneous photometric stereo and masking. In *2019 International Conference on 3D Vision (3DV)*, pages 222–229. IEEE, 2019.
- [26] Bjoern Haefner, Zhenzhang Ye, Maolin Gao, Tao Wu, Yvain Quéau, and Daniel Cremers. Variational uncalibrated photometric stereo under general lighting. *arXiv preprint arXiv:1904.03942*, 2019.
- [27] Satoshi Ikehata. Cnn-ps: Cnn-based photometric stereo for general non-convex surfaces. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [28] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, 2017.
- [29] Wenzel Jakob. Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>.
- [30] Jiyoung Jung, Joon-Young Lee, and In So Kweon. One-day outdoor photometric stereo via skylight estimation. In *Pro-*

ceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 4521–4529, 2015.

[31] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.

[32] Ira Kemelmacher-Shlizerman and Ronen Basri. 3d face reconstruction from a single image using a single reference face shape. *IEEE transactions on pattern analysis and machine intelligence*, 33(2):394–405, 2010.

[33] W. Lai, J. Huang, N. Ahuja, and M. Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5835–5843, 2017.

[34] Alexandros Lattas, Stylianos Moschoglou, Baris Gecer, Stylianos Ploumpis, Vasileios Trianafyllou, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme: Realistically renderable 3d facial reconstruction "in-the-wild". In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[35] Junxuan Li, Antonio Robles-Kelly, Shaodi You, and Yasuyuki Matsushita. Learning to minify photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7568–7576, 2019.

[36] Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2475–2484, 2020.

[37] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. In *SIGGRAPH Asia 2018 Technical Papers*, page 269. ACM, 2018.

[38] Guosheng Lin, Fayao Liu, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for dense prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.

[39] G. Lin, A. Milan, C. Shen, and I. Reid. RefineNet: Multi-path refinement networks for high-resolution semantic segmentation. In *CVPR*, July 2017.

[40] Guilin Liu, Duygu Ceylan, Ersin Yumer, Jimei Yang, and Jyh-Ming Lien. Material editing using a physically based rendering network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2261–2269, 2017.

[41] Fotios Logothetis, Ignas Budvytis, Roberto Mecca, and Roberto Cipolla. Px-net: Simple, efficient pixel-wise training of photometric stereo networks. *arXiv preprint arXiv:2008.04933*, 2020.

[42] S. Lombardi and K. Nishino. Reflectance and natural illumination from a single image. In *Proceedings of European Conference on Computer Vision (ECCV)*, October 2012.

[43] S. Lombardi and K. Nishino. Reflectance and illumination recovery in the wild. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 38(1):129–141, January 2016.

[44] Feng Lu, Xiaowu Chen, Imari Sato, and Yoichi Sato. Symps: Brdf symmetry guided photometric stereo for shape and light source estimation. *IEEE transactions on pattern analysis and machine intelligence*, 40(1):221–234, 2017.

[45] Roberto Mecca and Yvain Quéau. Unifying diffuse and specular reflections for the photometric stereo problem. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9. IEEE, 2016.

[46] Roberto Mecca, Ariel Tankus, Aaron Wetzler, and Alfred M Bruckstein. A direct differential approach to photometric stereo with perspective viewing. *SIAM Journal on Imaging Sciences*, 7(2):579–612, 2014.

[47] Abhimitra Meka, Maxim Maximov, Michael Zollhoefer, Avishek Chatterjee, Hans-Peter Seidel, Christian Richardt, and Christian Theobalt. Lime: Live intrinsic material estimation. In *Proceedings of Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[48] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.

[49] Zhipeng Mo, Boxin Shi, Feng Lu, Sai-Kit Yeung, and Yasuyuki Matsushita. Uncalibrated photometric stereo under natural illumination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2936–2945, 2018.

[50] Takahiro Okabe, Imari Sato, and Yoichi Sato. Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions. In *2009 IEEE 12th International Conference on Computer Vision*, pages 1693–1700. IEEE, 2009.

[51] Yvain Quéau, François Lauze, and Jean-Denis Durou. Solving uncalibrated photometric stereo using total variation. *Journal of Mathematical Imaging and Vision*, 52(1):87–107, 2015.

[52] Yvain Quéau, Roberto Mecca, Jean-Denis Durou, and Xavier Descombes. Photometric stereo with only two images: A theoretical study and numerical resolution. *Image and Vision Computing*, 57:175–191, 2017.

[53] Yvain Quéau, Tao Wu, and Daniel Cremers. Semi-calibrated near-light photometric stereo. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 656–668. Springer, 2017.

[54] Yvain Quéau, Tao Wu, François Lauze, Jean-Denis Durou, and Daniel Cremers. A non-convex variational approach to photometric stereo under inaccurate lighting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 99–108, 2017.

[55] Konstantinos Rematas, Tobias Ritschel, Mario Fritz, Efstratios Gavves, and Tinne Tuytelaars. Deep reflectance maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4508–4516, 2016.

[56] Hiroaki Santo, Masaki Samejima, Yusuke Sugano, Boxin Shi, and Yasuyuki Matsushita. Deep photometric stereo network. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 501–509, 2017.

[57] Soumyadip Sengupta, Jinwei Gu, Kihwan Kim, Guilin Liu, David W. Jacobs, and Jan Kautz. Neural inverse rendering of an indoor scene from a single image. In *International Conference on Computer Vision (ICCV)*, 2019.

[58] Soumyadip Sengupta, Angjoo Kanazawa, Carlos D Castillo, and David W Jacobs. Sfsnet: Learning shape, reflectance and illuminance of faces in the wild'. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6296–6305, 2018.

[59] Soumyadip Sengupta, Hao Zhou, Walter Forkel, Ronen Basri, Tom Goldstein, and David Jacobs. Solving uncalibrated photometric stereo using fewer images by jointly optimizing low-rank matrix completion and integrability. *Journal of Mathematical Imaging and Vision*, 60(4):563–575, 2018.

[60] Boxin Shi, Zhe Wu Mo, Dinglong Duan, Sai-Kit Yeung, and Ping Tan. A benchmark dataset and evalution for non-lambertian and uncalibrated photometric stereo. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 41(2):271–284, 2019.

[61] Konstantin Sofiuk, Ilya Petrov, Olga Barinova, , and Anton Konushin. f-brs: Rethinking backpropagating refinement for interactive segmentation. *arXiv preprint arXiv:2001.10331*, 2020.

[62] Tatsumi Taniai and Takanori Maehara. Neural inverse rendering for general reflectance photometric stereo. *arXiv preprint arXiv:1802.10328*, 2018.

[63] X. Wang, D. F. Fouhey, and A. Gupta. Designing deep networks for surface normal estimation. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 539–547, 2015.

[64] Olivia Wiles and Andrew Zisserman. Silnet : Single- and multi-view reconstruction by learning from silhouettes. In *British Machine Vision Conference 2017, BMVC 2017, London, UK, September 4-7, 2017*. BMVA Press, 2017.

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

[66] Zexiang Xu, Kalyan Sunkavalli, Sunil Hadap, and Ravi Ramamoorthi. Deep image-based relighting from optimal sparse samples. *ACM Transactions on Graphics (TOG)*, 37(4):126, 2018.

[67] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. *European Conference on Computer Vision (ECCV)*, 2018.

[68] Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, and Long Quan. Recurrent mvsnet for high-resolution multi-view stereo depth inference. *Computer Vision and Pattern Recognition (CVPR)*, 2019.

[69] Ye Yu and William AP Smith. Inverserendernet: Learning single image inverse rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3155–3164, 2019.

[70] Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, and Mubarak Shah. Shape-from-shading: a survey. *IEEE transactions on pattern analysis and machine intelligence*, 21(8):690–706, 1999.

[71] Qian Zheng, Yiming Jia, Boxin Shi, Xudong Jiang, Ling-Yu Duan, and Alex C Kot. Spline-net: Sparse photometric stereo through lighting interpolation and normal estimation networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8549–8558, 2019.

[72] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networkss. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.