

DeepShadow: Neural Shape from Shadow

Asaf Karnieli¹, Ohad Fried¹, and Yacov Hel-Or¹

School of Computer Science, Reichman University
asafkarnieli@gmail.com
{ofried,toky}@idc.ac.il

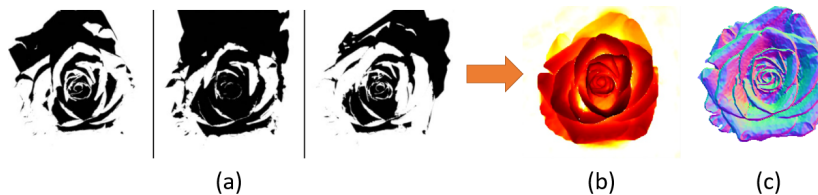


Fig. 1. Results on the rose object. (a) Input shadow maps, which were used for supervision (b) Depth produced by the algorithm. (c) Surface normals extracted from the depth map.

Abstract. This paper presents ‘DeepShadow’, a one-shot method for recovering the depth map and surface normals from photometric stereo shadow maps. Previous works that try to recover the surface normals from photometric stereo images treat cast shadows as a disturbance. We show that the self and cast shadows not only do not disturb 3D reconstruction, but can be used alone, as a strong learning signal, to recover the depth map and surface normals. We demonstrate that 3D reconstruction from shadows can even outperform shape-from-shading in certain cases. To the best of our knowledge, our method is the first to reconstruct 3D shape-from-shadows using neural networks. The method does not require any pre-training or expensive labeled data, and is optimized during inference time.

Keywords: Shape from Shadow, One-Shot, Inverse Graphics

1 Introduction

The photometric stereo setting was first defined in [25]. The setting includes a single camera viewpoint, and multiple varying illumination sources. Most works try to extract the underlying per-pixel normal map from the input images. The original problem assumed Lambertian objects, which only have a diffuse reflectance component. More recent works solve a more general setting, with various light and material properties, including specular components.

To date, most works extract the required information from the local illumination effect, defined by the bidirectional reflectance distribution function (BRDF), and either ignored the global cast shadow effect, or model it in a statistical way. Almost all recent works use Conv Nets to recover the 3D structure from photometric stereo. Since Conv Nets have limited receptive fields, global information can only be aggregated in the deeper layers where accurate spatial resolution is limited.

In our work, we extract the depth information directly from the cast shadows. As far as we know, this is the first attempt to do so using neural networks. In the photometric stereo setting, cast and attached shadow detection is a relatively easy learning task. The shadow maps are binary inputs, i.e., have a 0 or 1 value. To extract depth, we initially predict per-pixel depth information in 2D and thereafter produce a point-cloud in 3D from this prediction. The 3D points are then used to calculate global cast shadows, by tracing the light source to all destination pixels. The input images are used as supervision for the produced shadow maps. The entire process is differentiable, and the predicted depth can be learned using optimization. Our process is physics based and can aggregate global information while keeping the full resolution of the image.

Most previous works assume directional lights due to its simplicity, although this is usually not the case. Many scenes consist of point light sources, thus depth information should be extracted using a point light shading model. In this work, we assume images generated by point light sources, but we can easily extend our method to handle directional lights. Although the inputs in this work are only shadow maps, our method can be integrated with existing shape-from-shading models, in order to improve their results in complex shadowed scenes. Our code and data are available at <https://asafkar.github.io/deepshadow/>.

2 Related Work

Shape Reconstruction from Shadows Extracting shape information from shadows has been attempted in various works. This was mostly done before the deep learning era. An early attempt was performed in [8], which initializes a shadow map and sets lower and upper bounds on expected illuminated or shadowed pixels, which are then optimized until reaching a predicted height map. In [28], a graph-based representation is used to incorporate constraints induced by the shadow maps. This work also optimizes the height map using iteration based on low and high bounds. The ShadowCuts [2] work performs Lambertian photometric stereo with shadows present. It initially uses a graph-cut method for estimating light source visibility for every pixel. Then it uses the produced shadow maps along with shading information to perform surface normal integration, to produce the final predicted surface. In [26], the authors use 1D shadow graphs to speed up the 3D reconstruction process. 1D slices (i.e., rows or columns of the shadow images) are used, assuming the light source moves in two independent circular orbits. Each such slice can be solved independently, to retrieve the corresponding 1D slice of the underlying height map. This method

only handles light sources that are on a unit sphere, and the trajectories of the light sources must be perpendicular to each other. All the above works extract the shadow maps from the images using a hand-picked threshold. A recent method [16] uses an initial approximation of the 3D geometry as input to optimize object deformations by using soft shadow information. In contrary, our method does not require an initial approximation of the geometry. The method uses spherical harmonics to approximate the surface occlusion whereas we use implicit representations along with a linear-time tracing method.

Implicit Representations Implicit representations have recently been used to replace voxel grids or meshes with continuous parameterization using neural networks. These are used in [17] to query the color and opacity of a certain point in space, to produce the pixel color of an image acquired from a specific viewing position. Lately, implicit representations have also been used in works such as [1, 23, 30] to recover various underlying geometric and photometric properties such as albedo, normals, etc. These have all been done in multi-view settings and take many hours and days to optimize.

Viewshed Analysis Viewshed analysis solves the problem of which areas are visible from a specific viewpoint, and solved in [7, 10] using a line-of-sight (LOS) algorithm. This algorithm tracks along a ray from the viewpoint to the target point, while verifying that the target point is not occluded by the height points along the ray. Shapira [21] proposed the R3 method for viewshed computation, by generating LOS rays from the viewpoint to all other cells (pixels) in a height map. Franklin et al. [9] introduced the R2 algorithm, which optimizes the R3 method by launching LOS rays only to the outer cells (boundary pixels), while computing and storing the intermediate results of all cells on the way. This method is considered an approximation since some rays will encounter the same cells during the process.

Photometric Stereo In the past few years, various learning-based methods have been proposed to solve calibrated and uncalibrated photometric stereo problems [12, 5, 4, 27, 15]. Most existing methods solve the problem in a supervised manner, i.e., given inputs of N images and their associated light directions and intensities, the methods predict the surface normals as an output. The ground-truth normals are needed as supervision in order to train the network. In the uncalibrated scenarios, the light intensities and directions are not required. Works such as [5, 6] first regress the light direction and light intensities, and then solve the calibrated photometric stereo problem using the predicted lights. Methods such as [12, 24] solve the problem in an unsupervised manner, using a reconstruction loss. These methods only require the input images, which are used as self-supervision in the reconstruction loss.

Most of the above-mentioned methods use the pixelwise information to solve the problem. No inter-pixel relations are taken into account, besides the obvious local dependencies resulting from the local filtering. There are also methods such as [27, 15] that use local neighborhood pixels as well by feature aggregation. In [27], the authors use a graph-based neural network to unify all-pixel and per-pixel information, while in [15] the authors use each pixel as a data-point

to learn interactions between pixels. Although the latter method does handle cast shadows, it does so within a bound area of the observation map and in a statistical manner. Apart from this method, the previous methods ignore the cast shadows as a source of information, and treat them as a disturbance.

Some work solve the near-field photometric stereo problem, which assumes that the lights are close to the object, and thus cannot use directional lights. In [20], the authors use a hybrid approach that utilizes both distant and near-field models. The surface normals of a local area are computed by assuming they are illuminated by a distant light source, and the reconstruction error is based on the near field model. This method yields good results for close range lights, although it requires knowing the mean depth for initialization. It also requires the location of the point lights as well as the ambient lighting. The work in [14] also uses a far and near hybrid method, by first predicting the surface normals assuming far light sources, and then computing the depth by integrating the normal fields. The normal fields are then used to estimate light directions and attenuation that are used to compute the reflectance.

We also mention [19, 11], which produce high resolution depth maps by using low resolution depth maps and photometric stereo images. Since few photometric stereo-based methods target depth outputs, we will compare our work with [19].

In summary, most previous shape-from-shading methods need a large amount of data for training. Also, they ignore the cast shadow instead of using it as a source of information. The main contributions of our method are:

- We propose the first deep-learning based method for recovering depth and surface normals from shadow maps.
- Our method includes a shadow calculation component, which globally aggregates information in the spatial space.
- Our method uses linear-time calculation rather than quadratic complexity in [17] and its follow-up works. We also use 2D parameterization for depth maps, avoiding the expensive 3D parameterization in NeRF like works.
- In contrast to most shape-from-shading methods, which use photometric stereo data, our method is insensitive to non-diffuse reflectance (specular highlights) or varying intensity lights, which also enables it to generate good results in both from the near-field and far-light photometric stereo settings.
- Lastly, our method estimates the depth-map in a one-shot manner, avoiding costly data collection and supervised training.

3 Shape-from-Shadow Method

DeepShadow is a technique that estimates the shape of an object or a scene from cast and attached shadows. The input data are multiple shadow maps and the location of their associated point light sources. In our work, these can later be relaxed to input images alone. Similarly to photometric stereo, the data are generated or captured with a single viewpoint and multiple illuminations. We assume each input shadow map is a byproduct of a single image taken under a single light source. In contrast to other methods that use directional lights,

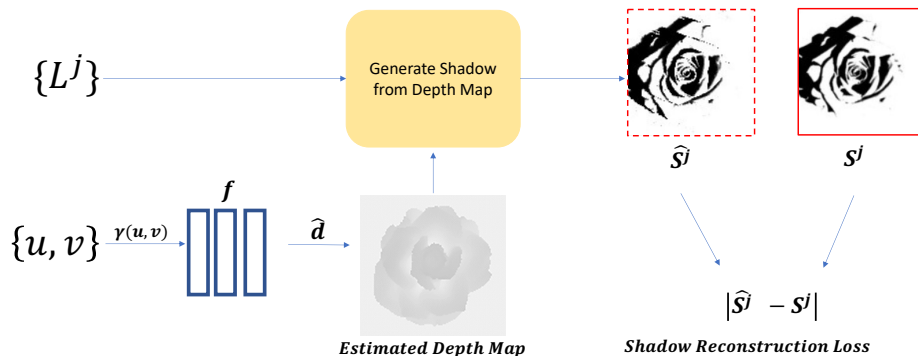


Fig. 2. Algorithm Overview. DeepShadow takes the light source location L^j and pixel coordinates (u, v) as inputs, along with the estimated depth \hat{d} from the MLP, and outputs an estimate of the shadow map \hat{S}^j at each pixel location. The ground-truth shadow map S^j is then used as a supervision to optimize the learned depth map.

we do not enforce the location of the point light to be on a unit sphere. Our algorithm can also work with directional light, which models a light source at infinity.

Fig. 2 shows our framework. First, a multi-layer perceptron (MLP) predicts the depth at each given location $\mathbf{u} = (u, v)$ in the image. Positional encoding $\gamma(\mathbf{u})$ is used at the input of the MLP. Then, the predicted depth \hat{d} along with an input light location L^j are used to estimate a shadow map \hat{S}^j . This is a physics based component which is differentiable but not trainable. The associated ground-truth shadow map S^j is used for supervision to optimize the MLP.

3.1 Shadow Map Estimation

To calculate the shadow of a pixel with respect to a specific light source L^j , the ‘shadow line scan’ algorithm is used. This algorithm traces a light ray from the light source to the destination pixel, and determines whether the pixel is illuminated or shadowed. More details are provided in Section 3.2.

To produce a shadow map from an estimated depth map, we use both world coordinates¹ $\mathbf{X} = \{x, y, z\}$ and image coordinates¹ $\mathbf{u} = \{u, v\}$ on the image plane. We assume the calibrated camera model is known. During the process, we use perspective projection and unprojection to go from one coordinate system to the other. Thus for a certain pixel location \mathbf{u}_i , $\mathbf{u}_i = P\mathbf{X}_i$, where P is the projection matrix $P = K \cdot [R|\mathbf{t}]$, K is the intrinsic matrix and $[R|\mathbf{t}]$ is the extrinsic matrix.

For a given light source $L^j = (L_x^j, L_y^j, L_z^j)$ in world coordinates, we initially map the light source to a point $\ell^j = P \cdot L^j$ on the image plane. To estimate the

¹ We omit the homogeneous coordinates for the sake of clarity.

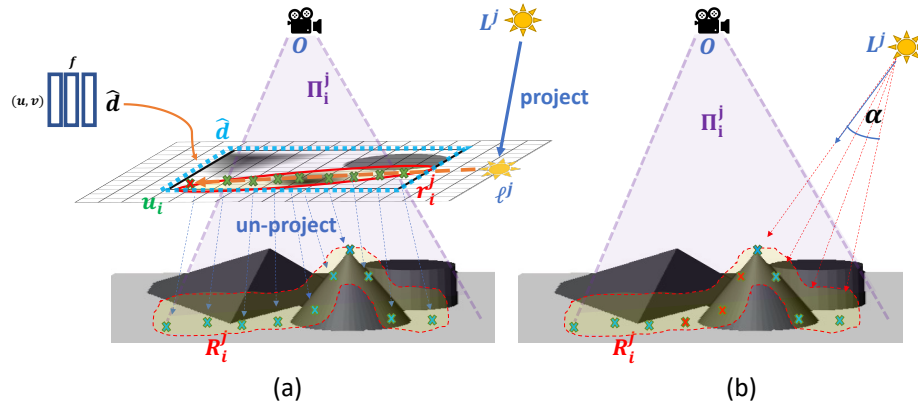


Fig. 3. Flow of shape-from-shadow. (a) The light source L^j is projected onto the image plane to receive ℓ^j . A ray \mathbf{r}_i^j of (u, v) points is created between ℓ^j and \mathbf{u}_i . Then, each point with its estimated depth \hat{d} is unprojected to world coordinates. (b) The shadow line scan algorithm is used on points in 3D space to calculate shadowed pixels. Red points are shadowed, since their angle to the light source is smaller than α .

shadow for a chosen pixel $\mathbf{u}_i = (u_i, v_i)$ in the image, a line \mathbf{r}_i^j of points in the image plane is generated between ℓ^j and \mathbf{u}_i :

$$\mathbf{r}_i^j(\alpha) = (1 - \alpha)\ell^j + \alpha\mathbf{u}_i, \quad \alpha \in [0, 1] \quad (1)$$

Points on \mathbf{r}_i^j that are outside the image frame are excluded. For each pixel in $\mathbf{r}_i^j(\alpha)$, we estimate the depth $\hat{d}_i^j(\alpha)$ by querying the MLP model at the pixel's coordinates. Each such triplet $(\mathbf{r}_i^j(\alpha)[x], \mathbf{r}_i^j(\alpha)[y], \hat{d}_i^j(\alpha))$ is then unprojected to world coordinates to receive its 3D location $R_i^j(\alpha)$ as described in Equation (2) and illustrated in Fig. 3a.

$$R_i^j(\alpha) = P^{-1}\mathbf{r}_i^j(\alpha) \cdot \hat{d}_i^j(\alpha), \quad \alpha \in [0, 1] \quad (2)$$

Once the 3D coordinates for each pixel in R_i^j are obtained, we can solve this as a 1D line-of-sight problem, as illustrated in Fig. 3b. Note, that this process will determine, *for each* point in R_i^j , whether it is shadowed or illuminated (and not only for the point \mathbf{u}_i). We use the image plane coordinates for parameterization of the depth, thus avoiding a costly full 3D parameterization (as in [17] and follow up works). This process assumes the object's depth is a function of (u, v) , which is a justifiable assumption since we are solving for the case of a single viewpoint.

Note that the light source L^j , viewer location O and a given point \mathbf{u}_i create a plane Π_i^j in 3D, as can be seen in Fig. 3b. ℓ^j is also located on Π_i^j , as are all the points in \mathbf{r}_i^j and all the points in R_i^j . This enables us to use a 1D line-of-sight algorithm between L^j and all points in R_i^j (elaborated in Section 3.2).

We emphasize that the shadow calculation depends only on the locations of the light source and the depth map. The camera’s center of projection (COP) O is used to generate the scan order for points along a 1D ray, and to expedite calculations as will be described in the next section.

3.2 Shadow Line Scan Algorithm

Given a height map and a viewer location, one can analyze all points visible to the viewer using a ‘line-of-sight’ calculation. This process identifies which areas are visible from a given point. It can be naively achieved by sending 1D rays to every direction from the viewer, and calculating the line-of-sight visibility for each ray. This is analogous to producing shade maps – the viewer is replaced with a light source L^j , and every pixel is analyzed to determine whether it is shadowed (visible from L^j) or not, in the same manner. We refer to this process as ‘shadow line scan’.

Calculating the visibility of each pixel point is time consuming and may take many hours to train. Instead, we propose to calculate the visibility of entire lines in the image. Each line R_i^j is calculated using a single scan. Each line includes the projections of all the points in 3D that are in the plane Π_i^j that is formed by the COP O , L^j and \mathbf{u}_i .

As is illustrated in Fig. 4, we define a vector from the light source L^j to each point in R_i^j , as

$$V_i^j(\alpha) = L^j - R_i^j(a). \quad (3)$$

Since we use the shadow maps as supervision and the shadow maps are composed of discrete pixels, we use discrete alpha values $\alpha \rightarrow \{\alpha_i\}_{i=0}^T$. We calculate all the angles between $V_i^j[\alpha_0]$ and each of the points. The angle for the i^{th} point is defined as

$$ang[i] = \arccos\left(\frac{V_i^j[\alpha_0] \cdot V_i^j[\alpha_i]}{\|V_i^j[\alpha_0]\| \cdot \|V_i^j[\alpha_i]\| + \epsilon}\right) \quad (4)$$

For numerical stability, ϵ is added to the denominator. $ang[i]$ is then compared to all previous angles $\{ang[k] \mid k \leq i\}$. If the current angle is larger than all previous angles, the current point is visible from the light source location and thus has no shadow. Otherwise, the point (pixel) is shadowed. A visualization of the process can be seen in Fig. 4.

This process can be reduced to a cumulative maximum function as defined in Equation (5). To achieve the final shadow estimate in Equation (6) we use a Sigmoid function to keep the final values between 0 and 1.

$$s_L[\alpha_i] = \max(ang[\alpha_i], s_L[\alpha_{i-1}]) \quad (5)$$

$$s[\alpha_i] = 2\sigma(ang[\alpha_i] - s_L[\alpha_i]) \quad (6)$$

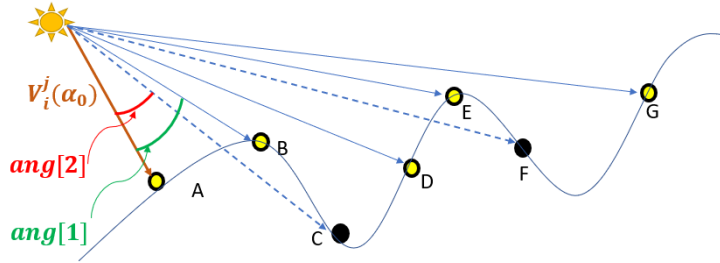


Fig. 4. 1D Shadow Line Scan Algorithm for Shadow Calculation. Points A–B are visible to the viewer, thus have no shadows. The angle $ang[2]$ between A and C (red) is smaller than the angle $ang[1]$ between A and B (green), thus C is not visible. Similarly, the angle between A and F is smaller than the last encountered largest angle, A–E, making F shadowed as well.

The cumulative maximum function was chosen since it is differentiable, similar to the well-known ReLU function. The rest of the process is also differentiable, and can be used to learn the depth map similar to what is done in inverse rendering methods.

3.3 Learning Depth from Shadows

For a given light source L^j , a dense estimated shadow map is produced by generating the shadow prediction for each pixel. The shadow line scan algorithm is applied for many lines in the image, covering the entire set of pixels, in order to generate a dense shadow map \hat{S}^j . The ground-truth shadow map can then be used for supervision, to learn the depth. Once finished, the optimized MLP can be used to generate a dense predicted depth map.

3.4 Computational Aspects

In [17] and similar follow-up works, to predict the pixel color, an integral must be calculated along a ray of samples. This requires quadratic computation complexity relative to the image size. In our work, the calculation takes linear time, since shadow predictions for all pixels along a shadow line scan can be calculated in a cumulative manner. The intermediate shadow results for a particular pixel and the maximum angle thus far are stored and used in the next pixel calculation along the line. We used the boundary sampling method (R2) [9], which sends LOS rays from the light source to the boundary of image only, instead of sending rays to all pixels. Although this method is considered less accurate than the R3 [21] method which calculates results for all pixels, it requiring an order of magnitude fewer calculations. The R2 method can also be sub-sampled using a coarse-to-fine sampling scheme. In the early iterations, rather than taking all pixels in the boundary, one can take every k_{th} pixel. Since the light locations vary, we will most likely use all pixels in the depth map at least once.

3.5 Loss Function

We used a loss composed of the reconstruction loss and a depth-guided regularization loss. The general loss is:

$$\mathcal{L} = \frac{1}{N} \sum_j \mathcal{L}_{rec}^j + \lambda \mathcal{L}_d \quad (7)$$

where \mathcal{L}_{rec}^j is the reconstruction loss:

$$\mathcal{L}_{rec}^j = \frac{1}{HW} |S^j - \hat{S}^j|. \quad (8)$$

\mathcal{L}_d is a depth regularization term:

$$\mathcal{L}_d = \sum_{ij} \left| \partial_x \hat{d}_{ij} \right| e^{-\|\partial_x \bar{I}_{ij}\|} + \left| \partial_y \hat{d}_{ij} \right| e^{-\|\partial_y \bar{I}_{ij}\|}, \quad (9)$$

where $\hat{d}_{ij} = \hat{d}(u_i, v_j)$ is the estimated depth at pixel (u_i, v_j) , ∂_x and ∂_y are gradients in the horizontal and vertical directions, and \bar{I} is the average color over all input images. Similar to [29], we assume that the average image edges provide a signal for the object discontinuity. This regularization term helps the depth map to converge faster.

4 Experimental Results

We compare our method to several shape-from-shading methods². While these methods require illuminated images as input (and we only require binary shadow maps) we find that this comparison illustrates the complementary nature of our method to shape-from-shading. As we will show, results depend on the statistics of the input objects, and for several classes of objects we achieve superior results to shape-from-shading even though we use less data.

4.1 Implementation details

We represent the continuous 2D depth map as a function of $\mathbf{u} = (u, v)$. Similar to [17], we approximate this function with an MLP $F_w(\mathbf{u}) = \hat{d}$. We used six-layer MLP with a latent dimension of 128, along with *sine* activation functions, similar to [22]. The MLP was initialized according to [22].

We implemented this using PyTorch [18]. We optimized the neural network using the Adam optimizer [13] with an initial learning rate of 5×10^{-5} , which was decreased every 15 epochs by a factor of 0.9. We reduce the temperature of the Sigmoid function 3 times during the optimization process, using a specific

² Shape from *shadow* was previously studied (Section 2). However, all works precede the deep learning era and only target simple objects. We compare to shape from *shading* methods that can be applied to the complex shapes in our datasets.

schedule. We run the optimization until the loss plateaus, which typically takes an hour on an input of 16 images with spatial resolution of 256x256, using an Intel i9-10900X CPU and an NVIDIA GeForce RTX 2080 Ti. The current implementation uses a single-process and contains many occurrences of native Python indexing, thus is sub-optimal and can be further optimized.

Photometric-stereo data is not usually accompanied by labeled shadow masks. Thus, for datasets which have no ground-truth shadows or light directions, we implement a model to estimate these. We use the Blobby and Sculptures datasets from [3] along with our own shadow dataset in order to train the model. Details can be found in the supplementary material.

4.2 Datasets

We show results on nine synthetic objects: 3 from previous work [12] and 6 curated and rendered by us from Sketchfab³. The objects we selected span various shapes and sizes, and they also possess one key property: many 3D features that cast shadows. As we will show, when this property holds, our method outperforms shape-from-shading methods and thus complements them. Our objects were each rendered with 16 different illumination conditions. Scenes are illuminated by a single point-light source.

Since shadow maps are not available in the dataset of Kaya et al. [12], we have our previously mentioned shadow-extraction model to estimate these. For our rendered objects, we use the ground-truth shadow maps as inputs. Qualitative results on real objects (without ground-truth) are available in the supplementary material.

4.3 Comparisons

We compare to the state-of-the-art methods of Chen et al. [5], Santo et al. [20], and Peng et al. [19], that use photometric stereo inputs, with the latter requiring a down-sampled depth map as an additional input. We compare both depth maps [19, 20] and normals [5, 20] (according to the generated outputs of each method).

Fig. 5 shows a qualitative evaluation of depth map result quality. The method of Santo et al. produces a low frequency estimation of depth, but does not contain any high frequency details. The method of Peng et al. produces results which are closer to the ground truth, but still lacking in sharpness (e.g., in *Relief* and *Cactus*). Our method is overall closest to the ground truth depth. Quantitatively, we measure normalized mean depth error (nMZE)⁴ in Table 1. We achieve lower nMZE on 5 out of the 6 objects, with an overall average nMZE which is 1.7 times [19] and 4.9 times [20] lower than the alternative.

Fig. 6 shows a qualitative evaluation of surface normal result quality. The method of Santo et al. produces overly smooth normals, with some examples

³ <https://sketchfab.com/>

⁴ we normalize each depth map by its own std and bias, in order to make the comparisons scale and bias invariant

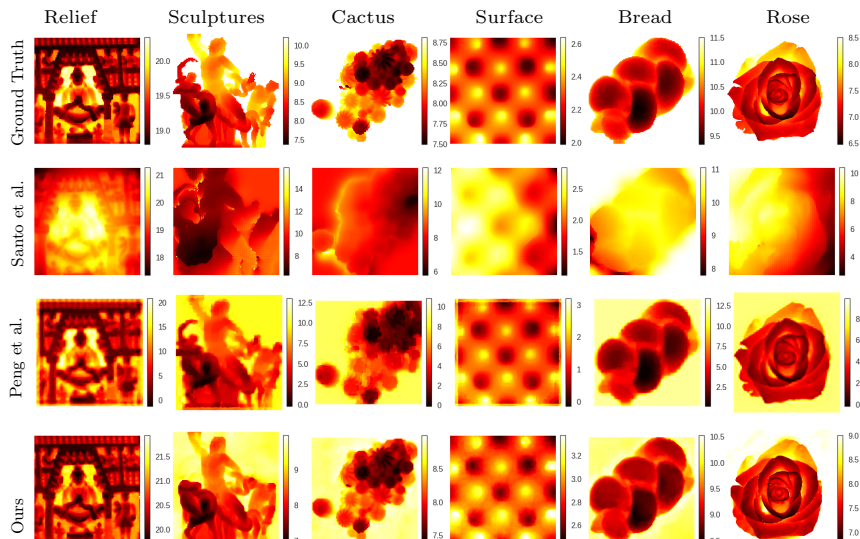


Fig. 5. Comparison of depth maps. Each row refers to a specific method.

relatively close to the ground truth (e.g., *Relief*) but others farther away (e.g., *Rose*). The method of Chen et al. is comparable to ours, with generally good results that resemble the ground truth normals, but lack sharp transitions. Quantitatively, we measure mean angle error (MAE in degrees) in Table 1. Our results are comparable to (albeit slightly better than) those of Chen et al., and 1.6 times better than those of Santo et al.

Table 1. Quantitative evaluation. The top three rows are depth map results, and the bottom rows are surface normal results. All methods but Peng et al. were given ground-truth light location, and Peng et al. was given a down-scaled depth map.

Method	Metric	Cactus	Rose	Bread	Sculptures	Surface	Relief	Avg
Santo et al. [20]	nMZE	0.96	1.16	0.77	0.81	0.75	0.75	0.87
Peng et al. [19]	nMZE	0.43	0.05	0.40	0.20	0.33	0.42	0.31
Chen et al. [5]	nMZE	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Ours	nMZE	0.33	0.11	0.16	0.19	0.10	0.18	0.18
Santo et al. [20]	MAE	32.79	50.23	33.95	54.49	22.21	21.93	35.93
Peng et al. [19]	MAE	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Chen et al. [5]	MAE	24.61	25.12	18.31	26.43	18.91	25.46	23.14
Ours	MAE	22.60	26.27	20.43	27.50	17.16	21.80	22.63

Note that [12] refers to a *rose* object as an object with ‘complex geometry and high amount of surface discontinuity’, where their method fails due to their

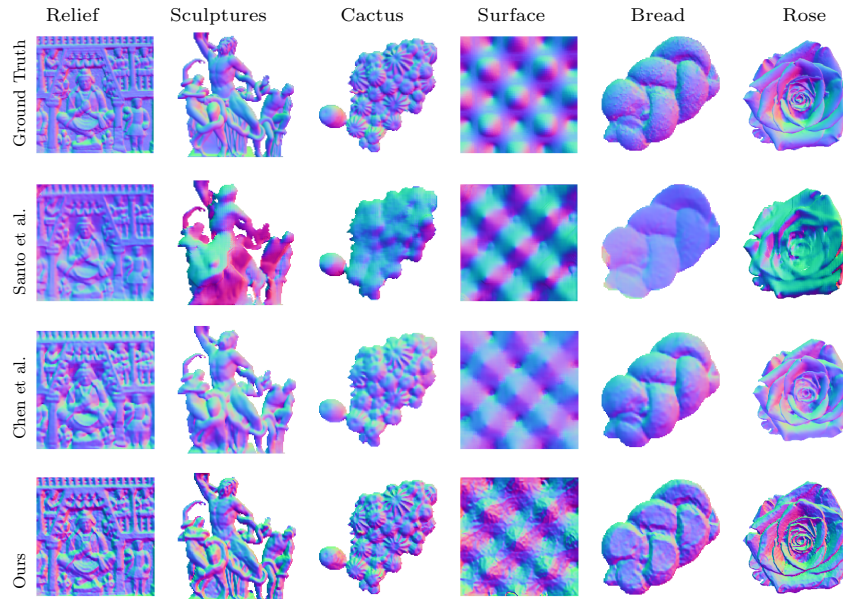


Fig. 6. Comparison of surface normals estimated by each of the compared methods. Each row represents a different method.

Table 2. Comparison of surface normal estimation on the dome dataset in [12] using calibrated methods (known light locations). Compared results were used from the original paper.

Method	Metric	Vase	Golf Ball	Face
Chen et al. [3]	MAE	27.11	15.99	16.17.81
Kaya et al. [12]	MAE	16.40	14.23	14.24
Ours	MAE	11.24	16.51	18.70

assumption of a continuous surface and due to a high amount of cast shadows. We render a similar object and show that our method does not require such assumptions. Visual results can be seen in Fig. 1.

Besides our own generated dataset, we compare our results to three different shape-from-shading results reported in [12]⁵ which can be seen in Table 2. Since the *vase* object is concave, it has a strong cast shadow affect, which is a strong learning signal. We can also see that the shape-from-shading methods fail to deal with this. The *golf* and *face* objects have sparse cast shadows, which is likely the reason why our method does not perform well there.

⁵ The code was not available at the time of writing this paper, so the reported results in [12] were used as-is.

Table 3. Average shadow pixels in image analysis. We measure MAE of estimated surface normals for increasing amount of average shadows. These are calculated by taking the average amount of shadowed pixels in each rendered object.

Avg. Shadow	A(48%)	B(60%)	C(66%)	D(69%)	E(60%)*	Avg
Chen et al. [5]	18.63	18.71	20.43	21.44	22.99	20.44
Santo et al. [20]	16.18	17.25	17.19	18.25	42.70	22.31
Ours	20.43	18.53	17.95	17.53	18.68	18.62

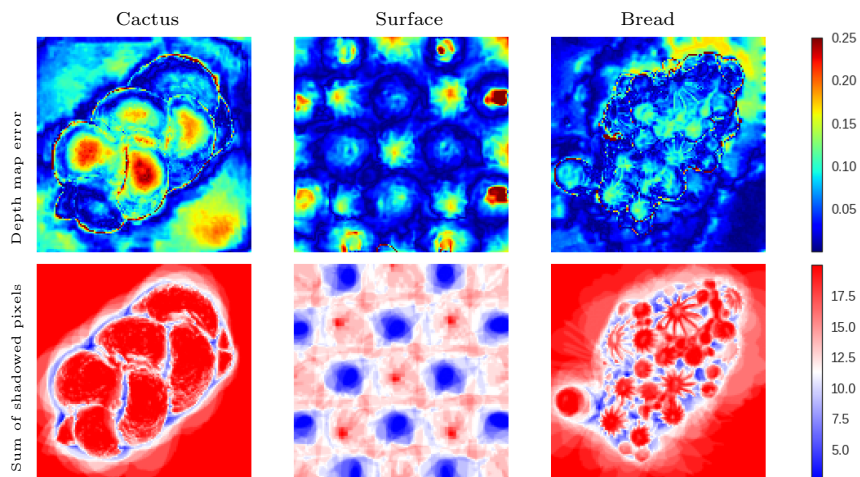


Fig. 7. Comparison of Depth Estimation Error (top) vs. Number of Shadowed Pixels (bottom). In the bottom row, we can observe pixels that are mostly shadowed (blue) or mostly illuminated (red) in all renderings. These pixels have higher depth estimation errors (top row). Pixels that are balanced, i.e., shadowed and illuminated in roughly equal proportions (white, bottom row) have a smaller depth estimation error.

4.4 Analysis & Ablation

Performance on various shadow amounts To further test our method, we generated the *surface* object using 16 lights placed at different incidence angles, thus creating inputs with increasing amount of shadows. We rendered this by placing the object at a constant location, and increasing the distance of the lights from the object at each different scene. We observe in Table 3 that DeepShadow’s accuracy improves as the amount of shadow present in the image increases, as opposed to the other tested methods, in which the accuracy degrades as the amount of shadowed pixels increases. Note that objects *A–D* were rendered with lights at a constant distance, and object *E* was rendered with lights at varying distances. We observe that [20] does not handle this scenario well.

Shadowed and illuminated pixels. We analyze the effects of the number of shadowed and illuminated pixels on the final depth estimation. In Fig. 7 we compare the error map (top) to the number of illuminated pixels across all

lighting conditions (bottom). Our method works best if this sum is balanced, i.e., if pixels are illuminated in some samples and shadowed in other samples. If pixels are always shadowed or always illuminated, reconstruction is ill-posed and our method does not perform well.

4.5 Failure Cases

Our algorithm is based on the line of sight algorithm, thus it relies on shadow gradients to reconstruct the underlying depth map. In objects such as the *face* in Table 2, the algorithm fails in flat areas, e.g., as the forehead, since the latter is relatively smooth and convex - and thus has no shadow from which to learn from. This is also true for the *bread object* in Table 1, which has large areas that are flat, as can be seen in Fig. 7. In contrast, the *vase* in Table 2 is relatively smooth yet also concave, and has many cast shadows, which enables our method to succeed.

5 Conclusions

In this work we proposed DeepShadow, a deep-learning based method for recovering shape from shadows. We show that in various scenarios, the shadow maps serve as good learning signals that enable the underlying depth and surface normals to be recovered. Experiments have shown that this method achieves results equal to or better than those generated by the various shape-from-shading algorithms, for severely shadowed objects, while using fewer or no data — since we are not using any training data besides the shadow maps. An additional benefit of our method is that we do not require knowing the light intensity, since the shadow maps are of binary values. Even though previous works have shown this can be estimated, the estimation may be erroneous, which affects the final result.

Our method fails on convex objects or objects with sparse shadows. In future work, this method should be combined with shape-from-shading methods in order for both methods to benefit from each other. Another possible research direction is producing super-resolution depth maps using the shadow clues, which would combine our work and [19]. Since we are using implicit representations, this may work well.

Acknowledgements This work was supported by the Israeli Ministry of Science and Technology under The National Foundation for Applied Science (MIA), and by the Israel Science Foundation (grant No. 1574/21).

References

- [1] Mark Boss et al. “NeRD: Neural Reflectance Decomposition from Image Collections”. In: *IEEE International Conference on Computer Vision (ICCV)*. 2021.

- [2] Manmohan Chandraker, Sameer Agarwal, and David Kriegman. “Shadow-Cuts: Photometric Stereo with Shadows”. In: *2007 IEEE Conference on Computer Vision and Pattern Recognition*. 2007, pp. 1–8. DOI: 10.1109/CVPR.2007.383288.
- [3] Guanying Chen, Kai Han, and Kwan-Yee K. Wong. “PS-FCN: A Flexible Learning Framework for Photometric Stereo”. In: *ECCV*. 2018.
- [4] Guanying Chen et al. *Deep Photometric Stereo for Non-Lambertian Surfaces*. 2020. arXiv: 2007.13145 [cs.CV].
- [5] Guanying Chen et al. *Self-calibrating Deep Photometric Stereo Networks*. 2019. arXiv: 1903.07366 [cs.CV].
- [6] Guanying Chen et al. “What Is Learned in Deep Uncalibrated Photometric Stereo?” In: *Computer Vision – ECCV 2020*. Ed. by Andrea Vedaldi et al. Cham: Springer International Publishing, 2020, pp. 745–762.
- [7] Richard Cole and Micha Sharir. “Visibility problems for polyhedral terrains”. In: *Journal of symbolic Computation* 7.1 (1989), pp. 11–30.
- [8] M. Daum and G. Dudek. “On 3-D surface reconstruction using shape from shadows”. In: *Proceedings. 1998 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No.98CB36231)*. 1998, pp. 461–468. DOI: 10.1109/CVPR.1998.698646.
- [9] Wm Randolph Franklin and Clark Ray. “Higher isn’t necessarily better: Visibility algorithms and experiments”. In: *Advances in GIS research: sixth international symposium on spatial data handling*. Vol. 2. Taylor & Francis Edinburgh. 1994, pp. 751–770.
- [10] Michael F Goodchild and Jay Lee. “Coverage problems and visibility regions on topographic surfaces”. In: *Annals of Operations Research* 18.1 (1989), pp. 175–186.
- [11] Bjoern Haefner et al. “Photometric Depth Super-Resolution”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2019.
- [12] Berk Kaya et al. *Uncalibrated Neural Inverse Rendering for Photometric Stereo of General Surfaces*. 2021. arXiv: 2012.06777 [cs.CV].
- [13] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: 1412.6980 [cs.LG].
- [14] Fotios Logothetis et al. *A CNN Based Approach for the Near-Field Photometric Stereo Problem*. 2020. arXiv: 2009.05792 [cs.CV].
- [15] Fotios Logothetis et al. *PX-NET: Simple and Efficient Pixel-Wise Training of Photometric Stereo Networks*. 2021. arXiv: 2008.04933 [cs.CV].
- [16] Linjie Lyu et al. “Efficient and Differentiable Shadow Computation for Inverse Problems”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. Oct. 2021, pp. 13107–13116.
- [17] Ben Mildenhall et al. “NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis”. In: *ECCV*. 2020.
- [18] Adam Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach et al. Curran Associates, Inc., 2019, pp. 8024–8035. URL: <http://papers.neurips.cc/paper/9015-pytorch->

- an-imperative-style-high-performance-deep-learning-library.pdf.
- [19] Songyou Peng et al. “Depth Super-Resolution Meets Uncalibrated Photometric Stereo”. In: *IEEE International Conference on Computer Vision (ICCV) Workshop*. 2017.
 - [20] Hiroaki Santo, Michael Waechter, and Yasuyuki Matsushita. “Deep Near-Light Photometric Stereo for Spatially Varying Reflectances”. In: *Computer Vision – ECCV 2020*. Ed. by Andrea Vedaldi et al. Cham: Springer International Publishing, 2020, pp. 137–152. ISBN: 978-3-030-58598-3.
 - [21] Andrew Shapira. “Visibility and terrain labeling”. PhD thesis. Rensselaer Polytechnic Institute, May, 1990.
 - [22] Vincent Sitzmann et al. “Implicit Neural Representations with Periodic Activation Functions”. In: *Proc. NeurIPS*. 2020.
 - [23] Pratul P. Srinivasan et al. *NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis*. 2020. arXiv: 2012.03927 [cs.CV].
 - [24] Tatsunori Tani and Takanori Maehara. *Neural Inverse Rendering for General Reflectance Photometric Stereo*. 2018. arXiv: 1802.10328 [cs.CV].
 - [25] Robert J. Woodham. “Photometric Method For Determining Surface Orientation From Multiple Images”. In: *Optical Engineering* 19.1 (1980), pp. 139–144. DOI: 10.1117/12.7972479. URL: <https://doi.org/10.1117/12.7972479>.
 - [26] Yukihiro Yamashita, Fumihiko Sakaue, and Jun Sato. “Recovering 3D Shape and Light Source Positions from Non-planar Shadows”. In: *2010 20th International Conference on Pattern Recognition*. 2010, pp. 1775–1778. DOI: 10.1109/ICPR.2010.1153.
 - [27] Zhuokun Yao et al. “Gps-net: Graph-based photometric stereo network”. In: *Advances in Neural Information Processing Systems* 33 (2020).
 - [28] Yizhou Yu and Johnny Chang. “Shadow Graphs and 3D Texture Reconstruction”. In: *International Journal of Computer Vision* 62 (Apr. 2005), pp. 35–60. DOI: 10.1023/B:VISI.0000046588.02227.3b.
 - [29] Huangying Zhan et al. “Self-supervised Learning for Single View Depth and Surface Normal Estimation”. In: *CoRR* abs/1903.00112 (2019). arXiv: 1903.00112. URL: <http://arxiv.org/abs/1903.00112>.
 - [30] Xiuming Zhang et al. “NeRFactor”. In: *ACM Transactions on Graphics* 40.6 (Dec. 2021), pp. 1–18. ISSN: 1557-7368. DOI: 10.1145/3478513.3480496. URL: <http://dx.doi.org/10.1145/3478513.3480496>.