

Neural Radiance Fields Approach to Deep Multi-View Photometric Stereo

Berk Kaya¹ Suryansh Kumar¹ Francesco Sarno¹ Vittorio Ferrari² Luc Van Gool^{1,3}
Computer Vision Lab, ETH Zürich¹, Google Research², KU Leuven³

Abstract

We present a modern solution to the multi-view photometric stereo problem (MVPS). Our work suitably exploits the image formation model in a MVPS experimental setup to recover the dense 3D reconstruction of an object from images. We procure the surface orientation using a photometric stereo (PS) image formation model and blend it with a multi-view neural radiance field representation to recover the object's surface geometry. Contrary to the previous multi-staged framework to MVPS, where the position, iso-depth contours, or orientation measurements are estimated independently and then fused later, our method is simple to implement and realize. Our method performs neural rendering of multi-view images while utilizing surface normals estimated by a deep photometric stereo network. We render the MVPS images by considering the object's surface normals for each 3D sample point along the viewing direction rather than explicitly using the density gradient in the volume space via 3D occupancy information. We optimize the proposed neural radiance field representation for the MVPS setup efficiently using a fully connected deep network to recover the 3D geometry of an object. Extensive evaluation on the DiLiGenT-MV benchmark dataset shows that our method performs better than the approaches that perform only PS or only multi-view stereo (MVS) and provides comparable results against the state-of-the-art multi-stage fusion methods.

1. Introduction

Estimating the 3D geometry of an object from images has been a fundamental research problem in computer vision for several decades [6, 44, 45]. A reliable solution has many real-world applications such as shape analysis, shape manipulation, molding, image rendering, forensics, *etc.* Even though there exist several ways to solve this problem, multi-view photometric stereo is one of the popular setups to recover the 3D geometry of the object (see Fig.1).

In the past, many active and passive 3D reconstruction approaches or pipelines were proposed to solve 3D reconstruction of objects [78, 100, 24, 129, 75, 56, 55]. However,

when it comes to the accuracy of recovered 3D shapes for its use in scientific and engineering purposes (metrology), methods that use only MVS or PS suffer [76, 90, 25, 26]. As a result, a mixed experimental setup such as multi-view photometric stereo (MVPS) is generally employed [65]. In such a setup, complementary modalities are used to obtain better surface measurements, which are otherwise unavailable from an individual sensor or method. Accordingly, similar fusion-based strategies gained popularity for surface estimation [79, 129, 61, 82, 12]. One may also prefer to use two or more active sensors to receive the surface data estimates for fusion. Nevertheless, this paper focuses on the MVPS setup, where the subject is placed on a rotating base and for each rotation multiple images are captured using one LED light source at a time. The major motivation for such an approach is that the active range scanning strategies used for object's 3D acquisition, such as structured light [28, 122, 123], 3D laser scanners [21], RGB-D sensors [129] are either complex to calibrate or provide noisy measurements or both. Further, these measuring techniques generally provide incomplete range data with outliers that require serious efforts for refinement.

Among the passive methods [91, 107, 25, 102, 12, 57, 58, 59], Multi-View Stereo (MVS) is a key approach in the automated acquisition of 3D objects from a set of images taken from different viewpoints [25]. Under the assumption of scene rigidity (rigid Lambertian textured surfaces), traditional MVS uses feature correspondences across images to reconstruct dense object geometry [24]. Recently, neural view synthesis methods have shown great potential for this task. However, the quality of their 3D reconstruction is far from satisfactory and several follow-ups are trying to address them for different scenarios [76, 118, 117, 116, 13].

An alternative approach to recover surface details is photometric stereo (PS), a.k.a shape from shading [106, 5]. PS methods estimate the object's surface normals given multiple images illuminated by different light sources, but all captured from the same camera viewpoint [106]. It is excellent at recovering surface normals independent of depth estimates and works well for untextured objects, non-Lambertian surfaces with fine details [35, 108, 42, 41, 10, 14]. To estimate the overall shape of the object, we can

directly integrate the high-resolution normal map. However, the absence of constraints between surface elements and the lack of global geometric constraints fails to determine each of those surface components' relative positions correctly. Further, most photometric stereo methods assume isotropic material objects and may fail to handle objects with anisotropic material like a piece of wood [41, 101, 95, 42].

Hence, MVS and PS complementary behavior in surface reconstruction from images helps us efficiently recover object shape. Using MVPS images, we utilize a CNN-based PS method to estimate surface normals using multiple light sources from a fixed viewpoint. For each view, multiple light sources are used to estimate the surface normal. We blend those surface normal estimates with a neural radiance field representation of the scene to recover the 3D reconstruction of the object. Existing state-of-the-art methods to this problem generally apply a sequence of steps [82, 65, 88] (i) procure the 3D position measurement using multi-view images or 3D scanner (ii) estimate surface orientation or iso-depth contours using photometric stereo methods, and (iii) fuse the surface orientation and 3D position estimates to recover 3D geometry using appropriate mathematical optimization. Now, several fusion strategies exist that can combine these alternate sources of information for better 3D shape reconstruction. [82, 79, 8, 83, 65, 88]. Of course, the precise steps taken by such approaches can provide better results, yet they rely heavily on explicit mathematical modeling [8, 82, 12, 119, 34, 83, 65], and **complex multi-staged network design** [88] which are complicated to execute. In contrast, our work provides a **simple** and general method that can suitably introduce the surface details coming from PS to neural radiance field representation of the scene for favorable performance gain.

Our work is inspired by the idea of local region-based techniques for volumetric illumination, which can render more realistic images [62, 48]. The ray-traced volume rendering approximates the surface normals using the gradient of the density along each (x, y, z) direction in volumetric space [84, 48]. However, to use such a notion for our problem setup, we must know the occupancy of the point in the volume space. To keep it simple, we utilize the surface normal from PS as the gradient density information for each sample point along the ray to reconstruct MVPS images as close as possible and recover 3D geometry. While one could recover depth from surface normals and then infer the occupancy of the volume density, we know that normal integration may lead to inaccurate depth estimates, hence incorrect occupancy information. So, we adhere to the proposed idea of using local gradients and show the validity of our method via experimental evaluations.

Our approach first estimates the surface normal for all the views using a deep photometric stereo network which




Figure 1. General setup for multi-view photometric stereo. The object is placed on a base with fixed rotation. A camera is placed at the center of the disk to capture images. LEDs are placed in a concentric ring for controlled light setup [65].

is trained independently in a supervised setting. We use sample spatial location x_i , the viewing direction d , and the object's surface normal for our MVPS representation. Surface normals for each 3D sample point along all known view directions are employed. For 3D reconstruction, our method optimizes a multi-layer perceptron that regresses from the position, view direction, and surface normal to single volume density and view-consistent RGB color. We use Fourier feature encoding on both positional and surface orientation data before passing them to the network [86]. Experimental results show that our method provides comparable or better results than previously proposed complex multi-view photometric stereo methods [8, 82, 12, 119, 34, 79], stand-alone state-of-the-art multi-view method [105], and view synthesis approach [76]. We evaluate the results on the DiLiGenT MVPS benchmark dataset [65, 94]¹. We make the following contributions:

- While previous multi-stage fusion methods to MVPS are very complex, we propose a much simpler continuous volumetric rendering approach that uses the local density gradient effects in MVPS image formation.
- Our work takes an opportunistic approach that exploits the complementary source of information in MVPS setup via two sets of representation *i.e.*, volumetric and surface.
- Despite being much simpler than fusion methods, our approach achieves better or comparable results than the state-of-the-art [83, 65], as well as stand-alone methods such as multi-view [105, 116], photometric stereo [41], and continuous volumetric rendering [73] methods.

2. Related Work

Here, we review important MVS, PS, view-synthesis and fusion based methods related to our work.

1. Multi-View Stereo (MVS) Methods. It aims at reconstructing a plausible 3D geometry of the object from a set

¹This dataset is publicly available for research purpose at: <https://sites.google.com/site/photometricstereodata/mv>

of images alone [24]. The working principle of MVS generally pivots around improving the dense image feature correspondence (local similarity) and camera parameters across images for triangulation [24, 60, 25, 7, 57, 58]. Recent developments in machine learning have led to the renovation of traditional MVS methods via deep-learning frameworks.

Roughly, it can be divided into four to five categories [24, 32]. (i) Volumetric methods require bounding box knowledge containing the subject. These methods estimate the relation between each voxel and the surface element, and their accuracy is greatly affected by voxel grid resolution [47, 51, 93, 23, 103]. (ii) Patch-based approach utilize the Barnes *et al.* [3] randomized correspondence idea in the scene space. Generally, these methods generate random 3D planes in the scene space and refine the depth representation and normal fields based on photo-consistency measures [7, 26, 25, 68]. (iii) Depth map reconstruction-based methods use a reference image with different source images under calibrated settings to improve the overall depth estimation [30, 99, 9, 92, 26, 110]. (iv) Point cloud-based methods operate on the 3D points and process the initial sparse point sets to densify the results [25, 63]. (v) Distributed structure-from-motion methods utilize the notion of motion averaging to improve large-scale 3D reconstruction [11, 125, 127]. Recently, deep neural networks have been widely adopted for MVS, which provide better performance than traditional MVS methods [32]. Earlier work in this area uses CNN’s for two-view [121] and multi-view stereo [33]. Lately, the learning-based MVS rely on the construction of 3D cost volume and use the deep neural networks for regularization and depth regression [18, 38, 113, 46, 71, 114, 111]. As most of these approaches utilize 3D CNN for cost volume regularization—which in general is computationally expensive, the majority of the recent work is motivated to meet the computational requirement with it. Few methods attempt to address it by down-sampling the input [114, 111]. Other attempts to improve the computational requirements uses sequential processing of cost volume [115], cascade of 3D cost volumes [19, 32, 112, 105], small cost volume with point-based refinement [18], sparse cost volume with RGB and 2D CNN to densify the result [120], learning-based patch-wise matching [71, 105] with RGB guided depth map super-resolution [105].

2. Photometric Stereo (PS) Methods. It is an alternative approach to infer the surface geometry from light-varying images. Those images are captured from a camera placed at a fixed viewpoint [106]. Although PS requires a controlled experimental setup, it prevails in recovering surface details independent of the depth knowledge. Similar to MVS, advancements in deep learning have led to the development of deep PS methods [41, 15, 101, 14]. The primary motive of the deep PS method is to let the neural network learn the complicated surface reflectance from data, which otherwise

is a challenging modeling task for surfaces with unknown BRDF’s [101, 52].

Woodham’s introduced PS for Lambertian surfaces [106]. Since then, the majority of the work in PS pivot around improving it for non-Lambertian surfaces [29, 20, 31, 42, 108]. In general, traditional PS can be categorized into four groups. (i) Robust methods recover surface normals by assuming the simple diffuse reflectance property of the object. It treats non-Lambertian reflectance components to be local and sparse, hence treated as outliers. Accordingly, rank minimization [108], expectation-maximization [109], sparse Bayesian regression [43], RANSAC [80], and other outlier filtering mechanisms [77] are used to estimate surface normals. Yet, these methods cannot handle material with soft and broad specularity [94]. (ii) Analytic BRDF modeling methods explicitly model the object’s specularity as a sum of diffused and specular components [29, 20, 31]. Still, it applies only to a limited class of objects due to the hand-crafted modeling strategy. (iii) General reflectance modeling-based methods utilize the objects isotropic [95, 43] and anisotropic [36] properties to estimate unknown BRDF rather than explicitly modeling it. (iv) the example-based method uses known object’s material knowledge for reflectance modeling [97]. While some work also assumes that the object’s material can be expressed using a small number of basis material [35], others take advantage of virtual sphere rendered with different materials without using any physical reference object [40]. Recent PS research is aimed at learning the complex surface reflectance using data via deep neural networks. Early deep PS methods aimed at learning the mapping between surface normals and reflectance measurements in a calibrated setting *i.e.*, light sources direction is given at train and test time [89, 101, 41, 16, 126, 64]. Recently, uncalibrated deep PS methods have been proposed to overcome the challenges associated with light sources calibration [15, 14, 52] and resolve GBR ambiguity in photometric stereo [17].

3. View Synthesis for 3D reconstruction. In recent years, view synthesis methods, in particular, the Neural Radiance Fields (NeRF) method for scene representation, have proposed an interesting idea to recover the 3D geometry of the scene using multi-view images [76]. NeRF has generated a new wave of interest in 3D computer vision and has led to several follow-ups in 3D data acquisition from images [118, 73, 81, 85, 98, 117, 87, 13]. NeRF uses a fully-connected deep neural network to represent the scene geometry and radiance information implicitly. It renders photo-realistic views by implicitly encoding surface volume density and color via a multi-layer perceptron (MLP). Once MLP is trained, the 3D shape can be recovered by using the estimated volume density. One can model the implicit object’s surface normals and radiance field in a unified way but may need volume occupancy information. Further, to

couple the surface normals term with the surface 3D point coming from the same measurement source may not provide significant gain [116], especially when dealing with PS images, where shading plays a critical role.

4. Fusion based Methods. Although the developments in RGB-D and other portable active sensors have led to success with volumetric methods on shading-based refinement [8, 72, 128], our work relies on image data for this problem. Much of the existing work that utilizes the images for complementary measurements uses explicit mathematical modeling for precise 3D geometry [34, 83, 65, 69, 66, 88]. *On the contrary*, our work utilizes photometric stereo to estimate surface normals and then blends it with the shape volume density in neural radiance field representation to synthesize multi-view images and recover shapes. Table (1) summarizes some of the recent and early developments in the area of multi-view photometric stereo.

3. Proposed Approach

We denote $\mathcal{I}^v = \{I_1^v, \dots, I_{N_p}^v\}$ as the set of N_p PS images for a given view $v \in [1, N_m]$. Our method assumes a calibrated setting for MVPS, *i.e.*, all the light source directions and camera calibrations are known. The proposed approach considers the notable image formation models used in computer vision and computer graphics *i.e.*, photometric stereo image formation model [106] and the rendering equation [49]. These imaging models has its advantage depending on the experimental setup. Since we are solving the well-known MVPS problem, we can exploit the benefit of both photometric and multi-view stereo setup. For our work, we assume that the subject under study is a solid texture-less object, which is often studied in photometric stereo research [65]. Next, we describe the deep PS network, followed by our neural radiance field modeling for a multi-view photometric stereo setup.

Deep Photometric Stereo. For simplicity, let's assume a single view case for an object with reflective surface, whose appearance can be encoded by a bidirectional reflectance distribution function (BRDF) Φ_s with surface normals $\mathbf{N} \in \mathbb{R}^{3 \times p}$ [101, 52]. Here, p symbolizes the total number of pixels. When object's surface is illuminated by a point light source positioned in the direction $l_i \in \mathbb{R}^{3 \times 1}$, then the image $I_i \in \mathbb{R}^{c \times r}$ captured by a camera in the view direction $v \in \mathbb{R}^{3 \times 1}$ can be modeled as

$$I_i^v = e_i \cdot \Phi_s(\mathbf{N}, l_i, v) \cdot \max(\mathbf{N}^T l_i, 0) + \epsilon_i \quad (1)$$

where, $e_i \in \mathbb{R}_+$ denotes the intensity, (c, r) symbolizes the columns and rows of an image, ϵ_i is an additive error and $\max(\mathbf{N}^T l_i, 0)$ accounts for the attached shadows. Eq.(1) image formation model for photometric stereo has led to outstanding developments in recovering fine details of the surface [96]. Yet, modeling unknown reflectance properties of different objects remains a fundamental challenge.

Method	Base data	Shape representation	Optimization
Mostafa <i>et al.</i> [79]	Laser, PS images	Neural Network	Back-Propagation + EKf [104]
Nehab <i>et al.</i> [82]	Scanner [21], PS images	Mesh	Mesh (Linear)
Hernandez <i>et al.</i> [34]	MVPS images	Mesh	Mesh (Coupled)
Park <i>et al.</i> [83]	MVPS images	Mesh + 2D displacement	Parameterized 2D (Sparse Linear)
Li <i>et al.</i> [65]	MVPS images	Depth Contour + 3D points	Poisson surface [53] + [82]
Logothetis <i>et al.</i> [69]	MVPS images and SDF	Parameterized SDF	Variational Approach
Ours	MVPS images	Multi-layer Perceptron	Adam [54]

Table 1. Previous work on passive approach to 3D shape recovery using orientation and range estimates. Despite Mostafa *et al.* [79] and Nehab *et al.* [82] use active modality, we included them for completeness.

Consequently, we utilize deep neural networks to learn the complicated BRDF's from input data. We leverage the observation map based CNN model to estimate surface normal under a calibrated setting [41]. Unlike other supervised methods, it has rotational invariance property for isotropic material, handles unstructured images and lights well, and above all provides best performance known to us with acceptable inference time.

▷ *Observation map:* For each pixel, this map contains the normalized observed intensity values due to all the light sources. In a general PS setup, the light sources are located in a concentric way. Thus, a one-to-one mapping between the light source position $(l_x, l_y, l_z) \in \mathbb{R}^3$ and corresponding x-y coordinate projection $(l_x, l_y) \in \mathbb{R}^2$ is possible. Note that $l_x^2 + l_y^2 + l_z^2 = 1 \forall l_i$ *i.e.*, the unit vector in the direction of source. We construct the observation map $\Omega_j \in \mathbb{R}^{w \times w}$ for each pixel j using its intensity value across all the N_p images as outlined in **Algorithm 1**. Here, w is the size of the observation map and the function $\zeta : \mathbb{R} \mapsto \mathbb{Z}_{0+}$. The scalar η_j for j^{th} pixel is $\eta_j = \max(e_1/I_1^v(j), \dots, e_{N_p}/I_{N_p}^v(j))$ is the normalizing constant. Since the projected source vectors can take values from $[-1, 1]$, they are scaled appropriately to get positive integer values.

Algorithm 1 Observation Map Construction

```

for  $j \leftarrow \{1, \dots, p\}$ 
  for  $i \leftarrow \{1, \dots, N_p\}$ 
     $\Omega_j(\zeta(w \cdot \frac{(l_x^i+1)}{2}), \zeta(w \cdot \frac{(l_y^i+1)}{2})) = \eta_j \frac{I_i^v(j)}{e_i}$ 
  end
end
```

▷ *PS Neural Network Architecture Details:* Inspired by the DenseNet design [39], the PS network first performs a convolution with 16 output channels on the input observation maps. The other part of the network consists of two dense blocks, a transition layer, and two dense layers, which is then followed by a normalization layer to recover surface normals as output. The dense block is composed of one ReLU layer, one 3×3 convolutional layer and a 0.2 dropout layer. The transition layer is composed of a ReLU layer, 1×1 convolutional layer, 0.2 dropout layer, and an average pooling layer, which is placed in between the two dense blocks to modify the feature map size. We train PS network end-to-end separately in a supervised setting. The l_2 (MSE) loss between estimated and the ground-truth nor-

als is minimized using Adam optimizer [54] (see Fig. 2). **Neural Radiance Field Representation for Multi-View Photometric Stereo.** Recently, volume rendering techniques for view synthesis, in particular NeRF [76] has shown a great potential for 3D data-acquisition from multi-view images. It represent a continuous scene as a 5D vector-valued function *i.e.*, $\mathbf{x} = (x, y, z)$ for each 3D location and (θ, ϕ) for every 2D viewing direction. Given multi-view images with known camera pose, NeRF approximate the assumed continuous 5D scene representation with a MLP that maps the $(\mathbf{x}, \theta, \phi)$ to RGB color \mathbf{c} and volume density $\sigma \in \mathbb{R}_+$. Using the classical volume rendering work [50], it models the expected color $C(\mathbf{r})$ of the camera ray $r(t) = \mathbf{o} + t\mathbf{d}$ with near and far bound t_n, t_f as:

$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt, \quad (2)$$

$$\text{where } T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s))ds\right). \quad (3)$$

Here, \mathbf{d} is the unit viewing direction. $T(t)$ is the accumulated transmittance along ray from t_n to t which caters the notion that how much light is blocked earlier along the ray. Given N samples along the ray, the continuous integral in Eq:(2) is approximated using the quadrature rule [74]:

$$\tilde{C}(\mathbf{r}) \approx \sum_{i=1}^N T_i \alpha_i(\mathbf{x}_i) \mathbf{c}_i(\mathbf{x}_i, \mathbf{d}), \quad (4)$$

$$\text{where } \alpha_i(\mathbf{x}_i) = \left(1 - \exp(-\sigma(\mathbf{x}_i)\delta_i)\right), \quad (5)$$

$$\text{and } T_i = \prod_{j=1}^{i-1} (1 - \alpha_j). \quad (6)$$

Here, δ_i is the distance between the adjacent discrete samples and α_i encapsulate how much light is contributed by the ray i . By construction, Eq.(4) approximates the alpha composited color as a weighted combination of all sampled colors \mathbf{c}_i along the ray. For more details we refer the readers to Mildenhall *et al.* work [76]. Now, lets have a closer look at the following rendering equation [49, 116]:

$$L_o(\mathbf{x}_i, \omega_o) = L_e(\mathbf{x}_i, \omega_o) + \int_{\mathcal{S}} \Phi_r(\mathbf{x}_i, \mathbf{n}_i, \omega_j, \omega_o) L_j(\mathbf{x}_i, \omega_j) (\mathbf{n}_i \cdot \omega_j) d\omega_j. \quad (7)$$

The above equation suggests that the rendering of a surface element \mathbf{x}_i depends on the emitted light in the scene and the bidirectional reflectance distribution function (BRDF) Φ_r describing reflectance and the color property of the surface accumulated over the half-sphere \mathcal{S} centered at \mathbf{n}_i . The significance of including surface normal to Φ_r in the rendering equation is put forward by IDR work [116]. Here,

ω_o, ω_j is the outgoing light direction and negative direction of the incoming light, respectively. Φ_r accounts for proportion of light reflected from ω_j towards ω_o at \mathbf{x}_i . L_o the radiance directed outward along ω_o from a particular position \mathbf{x}_i , and L_e is the emitted radiance of the light. In general, the light does not always hit the surface orthogonally, and so the dot product between the \mathbf{n}_i and ω_j attenuates the incoming light at \mathbf{x}_i . Hence, by restricting Φ_r and light radiance functions (L_o, L_e), which the radiance fields approximation can represent, we condition the color function to include the notion of density gradient for image rendering. Consequently, we use the normal estimated from PS to condition \mathbf{c}_i in Eq.(4). We rely on a deep photometric stereo network to estimate surface normals, overcoming BRDF modeling complications and providing us an excellent surface detail. Hence, our method has the inherent benefit over an entangled surface normal representation in image rendering [116]. Accordingly, we modify the Eq.(4) as follows:

$$\tilde{C}(\mathbf{r}) \approx \sum_{i=1}^N T_i \alpha_i(\mathbf{x}_i) \mathbf{c}_i(\mathbf{x}_i, \mathbf{n}_i^{ps}, \mathbf{d}). \quad (8)$$

Further, adding image features in Eq.(8) could be advantageous as in [118, 13]. However, MVPS setup generally deals with non-textured surfaces where using image features is not much of help. Still, an obvious advantage with the setup is that better surface details can be captured from shading. For simplicity, we use surface normals to condition volume rendering and refrain from relying on image features. So, our approach blends density gradient information into the continuous volume rendering formulation bypassing the explicit volume occupancy information. Concretely, we feed surface normals for each 3D sample point along the viewing direction to the neural rendering network.

Optimization and Loss Function. Following neural radiance fields optimization strategy [76], we encode each sampled position \mathbf{x}_i along the ray, viewing direction \mathbf{d} , and photometric stereo surface normal \mathbf{n}_i^{ps} using the Fourier features $\gamma(\mathbf{x}) = [\sin(\mathbf{x}), \cos(\mathbf{x}), \dots, \sin(2^{L-1}\mathbf{x}), \cos(2^{L-1}\mathbf{x})]$. We used $L = 10$ for $\gamma(\mathbf{x}_i)$, $L = 4$ for $\gamma(\mathbf{n}_i^{ps})$ and $L = 4$ for $\gamma(\mathbf{d})$. For efficient estimation of continuous integral (Eq.(8)) using quadrature rule, we use the stratified sampling approach to partition the near and far bound $[t_n, t_f]$ into N evenly-spaced discrete samples [74].

$$t_i \sim \mathcal{U}\left[t_n + \frac{i-1}{N}(t_f - t_n), t_n + \frac{i}{N}(t_f - t_n)\right]. \quad (9)$$

We employed MLPs (Multi-layer Perceptron) to optimize the following loss function.

$$\mathcal{L}_{mvps} = \sum_{\mathbf{r} \in \mathcal{B}} \|\tilde{C}_c(\mathbf{r}) - C(\mathbf{r})\|_2^2 + \|\tilde{C}_f(\mathbf{r}) - C(\mathbf{r})\|_2^2. \quad (10)$$

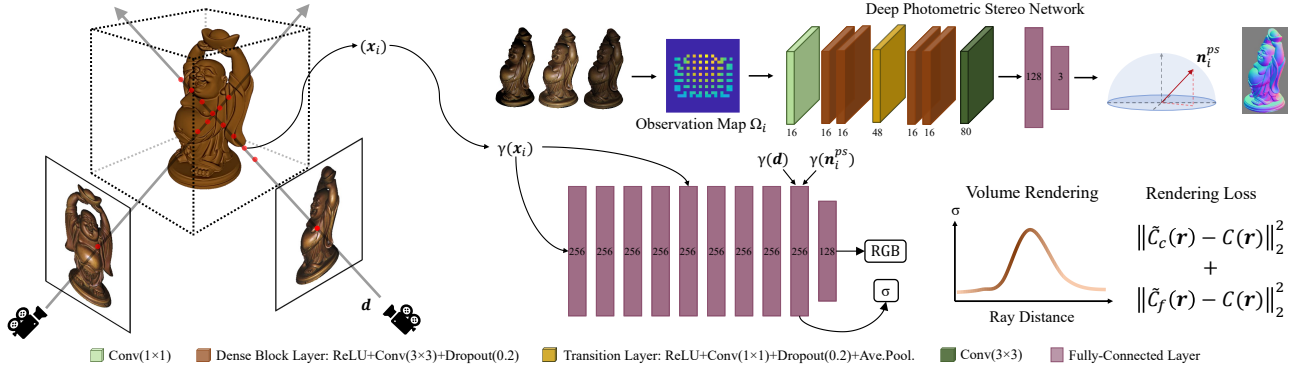


Figure 2. **Overview:** The deep photometric stereo network predicts surface normals of the object from each viewpoint using PS images. We model multi-view neural radiance fields by introducing gradient knowledge from PS network output in the density space for solving MVPS. Our work takes a much simpler approach than existing state-of-the-art multi-staged MVPS methods showing comparable accuracy.

where, \mathcal{B} denotes set of all rays in the batch. We used hierarchical volume sampling strategy to densely evaluate neural radiance field network for N query point along each ray. To that end, we first sample N_c points using the stratified sampling strategy and optimize the coarse network $\tilde{C}_c(\mathbf{r})$ (Eq.(8)). With the known output distribution of the coarse network, we sample N_f points using inverse transform sampling to optimize the fine network $\tilde{C}_f(\mathbf{r})$. In Eq.(10), the variable $C(\mathbf{r})$ is the observed color for the pixel (see Fig.2).

4. Experiment and Ablations

We conducted our experiments, statistical evaluations, and ablation study on the DiLiGenT-MV benchmark dataset [65, 94]. It consists of complex BRDF images taken from 20 viewpoints. For each viewpoint, 96 images are captured, each illuminated by a light source in a different known position (calibrated). The dataset includes 5 objects (BEAR, BUDDHA, COW, POT2, READING) with complex surface profiles and its images are captured under large illumination changes. For creating this dataset, the distance between the camera and object is set to 1500 mm [65].

Implementation Details. (a) *Deep-Photometric Stereo:* We train the deep photometric stereo network on the CyclesPS dataset [41]. We used the Adam optimizer [54] with a learning rate of 10^{-3} and trained for 10 epochs. A per-pixel observation map with a size of 32×32 is used at train and test time. During testing, we applied the network on 96 PS images per subject from DiLiGenT-MV dataset. (b) *MLP Optimization for MVPS:* For each object in DiLiGenT-MV dataset, we optimize a dense fully-connected neural network that is composed of 8 fully-connected ReLU layers. Each layer is composed of 256 channels. In addition to the density output from the 8th layer, 24 dimensional view-direction and surface normal Fourier features are then fed at the 9th layer for rendering (see Fig.2). We use $N_c = 64$ points for the coarse network and $N_f = 128$ points for the

fine network. We optimize the network parameters for 30 epochs with a batch size of 1024 rays and an initial learning rate of 10^{-4} . This takes 7 hours per object on an NVIDIA GeForce RTX 2080 Ti with 11GB RAM.

4.1. Baseline Comparison

For comparison against the baseline methods, we classified them into two categories: (a) *Standalone methods:* It either uses MVS or PS set up to recover the 3D shapes of an object. (b) *Multi-Staged Fusion methods:* It first recovers the sparse 3D point cloud of the object using multi-view images and surface orientation using fixed viewpoint PS images for each view. These spatial positions and orientations are then fused using a different method/pipeline to recover the 3D geometry. To compare against the standalone methods, we pick the 4th light source in DiLiGenT-MV setup. To recover the 3D shape using our method, we run our fine model by sampling points in 512^3 volumetric space uniformly. The recovered density is queried using marching-cube algorithm [70] with $\sigma = 10$. Table 2 compares the 3D reconstruction accuracy of our method to standalone and multi-stage fusion methods. We evaluate accuracy using the standard Chamfer-L1 distance metric between the recovered shape and the ground-truth shape after registration.

(a) Standalone Methods. For standalone baselines comparison, we compare our method with the recent state-of-the-art in MVS, PS, and View-Synthesis.

(1) *PatchMatch Network* [105]: This method has recently shown state-of-the-art performance in MVS. PatchMatch-Net proposed an end-to-end trainable network that has fast inference time and works well even for high-resolution images. It is trained on the DTU dataset [1]. We empirically observed that using two source frames per reference view on the DiLiGenT-MV dataset results in more accurate depth estimation. After getting the depth map from the network for each view, we transform the results to a point cloud by back-projecting the depth values to 3D space.

Method Type →	Standalone Methods (↓)				Multi-Stage Fusion Methods (↓)		
Dataset	PM-Net [105]	IDR [116]	NeRF [76]	Ours	R-MVPS [83]	B-MVPS [65]	Ours
BEAR	2.13	10.97	0.62	0.66	0.89	0.63	0.66
BUDDHA	0.72	5.18	0.99	1.00	0.64	0.40	1.00
COW	1.68	6.45	0.92	0.71	0.42	0.54	0.71
POT2	1.30	6.68	0.64	0.63	1.29	0.55	0.63
READING	1.64	7.05	1.22	0.82	0.98	0.85	0.82

Table 2. Quantitative 3D reconstruction accuracy comparison against the competing methods on DiLiGenT-MV benchmark. We used Chamfer-L1 metric to compute the accuracy. The statistics show that our method is better and comparable to the stand-alone methods and Multi-Stage fusion methods respectively. **Generally, Multi-Stage fusion methods are heuristic in nature and require careful execution at different stages.** On the contrary, our method is much simpler to realize and implement. Note: Even though we tried our best to make sure that IDR codes are compiled the way it’s on their paper code, we are surprised by IDR results and we are still investigating it. The above result is the best we achieved out of its publicly available code after several trials.

Method	BEAR	BUDDHA	COW	POT2	READING
CNN-PS [41]	0.78	0.83	0.87	0.86	0.89
Ours	0.09	0.11	0.11	0.07	0.06

Table 3. Multi-view depth error comparison against CNN-PS [41]. For comparison, we integrate the surface normals from CNN-PS to recover its depth which is scaled appropriately to compute the l_1 depth accuracy. For ours, we projected the recovered shape reconstruction to the estimated depth and corresponding depth accuracy.

(2) *IDR* [116]: This recently published method proposes a differentiable rendering formulation that is capable of implicitly modeling variety of lights and surface materials. IDR demonstrated state-of-the-art results on the DTU dataset and showed impressive 3D reconstructions from MVS images. For comparison, we train the model on DiLiGenT-MV for 2000 epochs. During training and testing, we maintained the default settings introduced by [116].

(3) *NeRF* [73]: Even though it is developed for novel view synthesis, recently NeRF has been widely used as a baseline for multi-view 3D reconstruction [118]. By sampling the volume density σ , it is possible to recover 3D geometry. We use an initial learning rate of 10^{-4} and batch size of 1024 rays. We run the optimization for 30 epochs and threshold the density values (512^3) at 10 to recover the 3D geometry.

(4) *CNN-PS* [41]: This method proposes a dense convolution neural network to learn the mapping between PS images and surface normals directly. It can handle non-convex surfaces and complicated BRDFs. We integrate the obtained surface normals using Horn and Brooks’ method [37] to get the depth map, and scale it to [-1, 1]. Next, we also project our 3D shape to the different cameras and recover the depth on a similar scale for statistical comparison. Table 3 shows our performance comparison against CNN-PS.

(b) Multi-Stage Fusion Methods. Fusion approaches to MVPS usually comprise several steps that are heuristic in nature, and proper care must be taken to execute all the steps well. For evaluation, we compared against the two well-known baselines in MVPS (see Table 2).

(1) *Robust MVPS* [83]. It employs a series of different algorithms to solve MVPS. It first uses multi-view images to recover the coarse 3D mesh of the object using structure from motion. Next, this mesh is projected to 2D planar space




Figure 3. Training and validation loss curve on BUDDHA seq.

for parameterization to estimate multi-view consistent surface normals using photometric stereo setup. Finally, the estimated surface normals are used for mesh refinement to recover the fine-detailed geometry of the object.

(2) *Benchmark MVPS* [65]. This method takes a set of steps to estimate a fine-detailed 3D reconstruction of an object from MVPS images. It first estimates iso-depth contours from the PS images [2] and then uses a structure-from-motion algorithm to recover a sparse point cloud from the MVS images. Later the depth of these 3D points is propagated along the iso-depth contour to recover the complete 3D shape. The spatially varying BRDF is computed once the 3D shape is recovered.

4.2. Analysis and Ablation

(a) *Training and Validation Analysis:* To demonstrate the effect of surface normal information on image rendering, we compare our method with NeRF. Fig.3 provides training and validation curves for both methods on BUDDHA. As expected, our method provides much higher image rendering quality during the learning process. Additionally, Table 4 provides PSNR and LPIPS [124] scores for IDR, NeRF and our method. These results show that our rendering quality is much better than other view-synthesis approaches.

(b) *Effect of Volume Sampling:* We test the importance of the sampling strategy by performing an ablation study on the BEAR dataset. In the first experiment, we reduced the number of coarse samples N_c to 16. This increased our reconstruction error from 0.66 to 0.84, indicating that an accurate coarse geometry is needed for better performance.

Dataset →	BEAR		BUDDHA		COW		POT2		READING	
Method	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓
IDR [116]	4.43	0.2370	9.87	0.2261	9.15	0.1571	7.71	0.1662	6.66	0.1815
NeRF [76]	29.97	0.0235	29.00	0.0455	30.80	0.0192	28.88	0.0269	28.12	0.0346
Ours	37.16	0.0122	33.59	0.0162	34.49	0.0134	30.47	0.0258	30.46	0.0311

Table 4. Quantitative image rendering quality comparison on DiLiGenT-MV benchmark. Our method can render much better images as it utilizes the surface normal information acquired with PS.

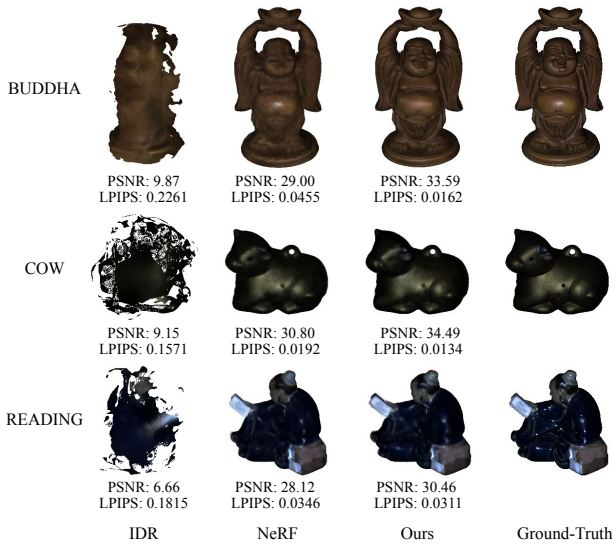


Figure 4. Comparison of volume rendering on DiLiGenT-MV achieved by IDR [116], NeRF [76] and our method. Without surface normals, NeRF lacks in details and produces blurry renderings. On the other hand, our method recovers fine surface details and renders accurate images by blending surface normal information in the volume rendering process. PSNR (higher the better), LPIPS (lower the better).

As a second experiment, we increased N_c to 256, but we did not use fine sampling. Here, the reconstruction error remained at 0.66, showing that more samples must get similar reconstruction quality.

(c) *Effects of multiple light sources on NeRF*: Since our method uses multiple light sources for normal estimation using deep PS, one could think about supplying multiple light source images to NeRF. So, rather than choosing the single best light source and apply NeRF to multi-view images, we test here what happens if we provide multiple light source images to it. The motivation of this experiment is to study the generalization of NeRF to (i) different light sources, however, for each experiment the light is consistent throughout the image sequence. (ii) what if the source is changing over the image sequence of the dataset. For this experiment, we sampled 20 lights out of 96 and experimented on the DiLiGenT-MV dataset. The minimal condition for PS to work is having more than three images [106] under Lambertian assumptions, and therefore, 20 light source is good enough. Table 5 shows the empirical results for setup (i). For more results, see supp. material.

Dataset	BEAR	BUDDHA	COW	POT2	READING
NeRF [76]	0.80	1.07	0.78	0.81	1.18
Ours	0.70	1.06	0.79	0.73	0.98

Table 5. Quantitative 3D reconstruction accuracy against NeRF [76]. We tested both approaches with 20 different light configurations. We provide average scores using Chamfer-L1 metric.

Limitations and Further Study. Our method combines two independent research fields that practice precise 3D reconstruction of an object from images. We demonstrated that the proposed method provides favorable results against the competing methods. However, we make a few assumptions, such as calibrated MVPS setup and solid objects, which limit our approach to broader adoption. Further, we do not explicitly model inter-reflections. Consequently, it would be interesting to extend our work to the uncalibrated settings. Also, exploring joint modeling of PS normal and surface normal from the implicit surface representation is left as a future extension. To improve the convergence time and modeling of uncertainty, we are in process of assessing [67, 4, 117, 27, 22].

5. Conclusion

We introduced a straightforward method that takes an opportunistic approach to solve the multi-view photometric stereo problem. We conclude that by exploiting the photometric stereo image formation model and the recent continuous volume rendering for multi-view image synthesis, we can reconstruct the 3D geometry of the object with high accuracy. Further, our formulation inherently utilizes the notion of the density gradient by leveraging the photometric stereo response and hence bypasses the explicit modeling of 3D occupancy information. Of course, using prior from structure-from-motion and its variants could be beneficial, as it could provide surface occupancy directly, but it would make the framework more complex. We demonstrated that introducing knowledge about the density gradient to the neural radiance field representation provides encouraging improvements in MVPS. We assessed the suitability of our work via extensive experiments on the DiLiGenT-MV benchmark dataset. We believe that our work can open up a new direction for study in multi-view photometric stereo for further developments.

Acknowledgement. This work was funded by Focused Research Award from Google (CVL, ETH 2019-HE-318, 2019-HE-323).

References

- [1] Henrik Aanæs, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders Bjarholm Dahl. Large-scale data for multiple-view stereopsis. *International Journal of Computer Vision*, 120(2):153–168, 2016.
- [2] Neil G Alldrin and David J Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.
- [3] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009.
- [4] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. *arXiv preprint arXiv:2103.13415*, 2021.
- [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] Andrew Blake and Andrew Zisserman. *Visual reconstruction*. MIT press, 1987.
- [7] Michael Bleyer, Christoph Rhemann, and Carsten Rother. Patchmatch stereo-stereo matching with slanted support windows. In *Bmvc*, volume 11, pages 1–11, 2011.
- [8] Erik Bylow, Robert Maier, Fredrik Kahl, and Carl Olsson. Combining depth fusion and photometric stereo for fine-detailed 3d models. In *Scandinavian Conference on Image Analysis*, pages 261–274. Springer, 2019.
- [9] Neill DF Campbell, George Vogiatzis, Carlos Hernández, and Roberto Cipolla. Using multiple hypotheses to improve depth-maps for multi-view stereo. In *European Conference on Computer Vision*, pages 766–779. Springer, 2008.
- [10] Manmohan Krishna Chandraker, Fredrik Kahl, and David J Kriegman. Reflections on the generalized bas-relief ambiguity. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 1, pages 788–795. IEEE, 2005.
- [11] Avishek Chatterjee and Venu Madhav Govindu. Efficient and robust large-scale rotation averaging. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 521–528, 2013.
- [12] Avishek Chatterjee and Venu Madhav Govindu. Photometric refinement of depth maps for multi-albedo objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 933–941, 2015.
- [13] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. *arXiv preprint arXiv:2103.15595*, 2021.
- [14] 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.
- [15] 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.
- [16] Guanying Chen, Kai Han, and Kwan-Yee K Wong. Psf-ncn: A flexible learning framework for photometric stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [17] 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.
- [18] Rui Chen, Songfang Han, Jing Xu, and Hao Su. Point-based multi-view stereo network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1538–1547, 2019.
- [19] Shuo Cheng, Zexiang Xu, Shilin Zhu, Zhuwen Li, Li Erran Li, Ravi Ramamoorthi, and Hao Su. Deep stereo using adaptive thin volume representation with uncertainty awareness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2524–2534, 2020.
- [20] Hin-Shun Chung and Jiaya Jia. Efficient photometric stereo on glossy surfaces with wide specular lobes. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [21] James Davis, Ravi Ramamoorthi, and Szymon Rusinkiewicz. Spacetime stereo: A unifying framework for depth from triangulation. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages II–359. IEEE, 2003.
- [22] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791*, 2021.
- [23] Olivier Faugeras and Renaud Keriven. *Variational principles, surface evolution, PDE’s, level set methods and the stereo problem*. IEEE, 2002.
- [24] Yasutaka Furukawa and Carlos Hernández. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015.
- [25] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2009.
- [26] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 873–881, 2015.
- [27] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. *arXiv preprint arXiv:2103.10380*, 2021.
- [28] Jason Geng. Structured-light 3d surface imaging: a tutorial. *Advances in Optics and Photonics*, 3(2):128–160, 2011.
- [29] Athinodoros S Georgiades. Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. In *ICCV*, pages 816–823. IEEE, 2003.

- [30] Michael Goesele, Brian Curless, and Steven M Seitz. Multi-view stereo revisited. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2402–2409. IEEE, 2006.
- [31] Dan B Goldman, Brian Curless, Aaron Hertzmann, and Steven M Seitz. Shape and spatially-varying brdfs from photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(6):1060–1071, 2009.
- [32] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2495–2504, 2020.
- [33] Wilfried Hartmann, Silvano Galliani, Michal Havlena, Luc Van Gool, and Konrad Schindler. Learned multi-patch similarity. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1586–1594, 2017.
- [34] Carlos Hernandez, George Vogiatzis, and Roberto Cipolla. Multiview photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3):548–554, 2008.
- [35] Aaron Hertzmann and Steven M Seitz. Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1254–1264, 2005.
- [36] Michael Holroyd, Jason Lawrence, Greg Humphreys, and Todd Zickler. A photometric approach for estimating normals and tangents. *ACM Transactions on Graphics (TOG)*, 27(5):1–9, 2008.
- [37] Berthold KP Horn and Michael J Brooks. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing*, 33(2):174–208, 1986.
- [38] Yuxin Hou, Juho Kannala, and Arno Solin. Multi-view stereo by temporal nonparametric fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2651–2660, 2019.
- [39] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [40] Zhuo Hui and Aswin C Sankaranarayanan. Shape and spatially-varying reflectance estimation from virtual exemplars. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(10):2060–2073, 2016.
- [41] 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.
- [42] Satoshi Ikehata and Kiyoharu Aizawa. Photometric stereo using constrained bivariate regression for general isotropic surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2179–2186, 2014.
- [43] Satoshi Ikehata, David Wipf, Yasuyuki Matsushita, and Kiyoharu Aizawa. Photometric stereo using sparse bayesian regression for general diffuse surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(9):1816–1831, 2014.
- [44] Katsushi Ikeuchi. Constructing a depth map from images. Technical report, MIT, Cambridge AI Lab, 1983.
- [45] Katsushi Ikeuchi. Determining a depth map using a dual photometric stereo. *The International journal of robotics research*, 6(1):15–31, 1987.
- [46] Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. Dpsnet: End-to-end deep plane sweep stereo. In *International Conference on Learning Representations*, 2018.
- [47] Mengqi Ji, Juergen Gall, Haitian Zheng, Yebin Liu, and Lu Fang. SurfacerNet: An end-to-end 3d neural network for multiview stereopsis. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2307–2315, 2017.
- [48] Daniel Jönsson, Erik Sundén, Anders Ynnerman, and Timo Ropinski. A survey of volumetric illumination techniques for interactive volume rendering. In *Computer Graphics Forum*, volume 33, pages 27–51. Wiley Online Library, 2014.
- [49] James T Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 143–150, 1986.
- [50] James T Kajiya and Brian P Von Herzen. Ray tracing volume densities. *ACM SIGGRAPH computer graphics*, 18(3):165–174, 1984.
- [51] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. *arXiv preprint arXiv:1708.05375*, 2017.
- [52] Berk Kaya, Suryansh Kumar, Carlos Oliveira, Vittorio Ferrari, and Luc Van Gool. Uncalibrated neural inverse rendering for photometric stereo of general surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3804–3814, 2021.
- [53] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006.
- [54] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [55] Suryansh Kumar. Jumping manifolds: Geometry aware dense non-rigid structure from motion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5346–5355, 2019.
- [56] Suryansh Kumar, Anoop Cherian, Yuchao Dai, and Hongdong Li. Scalable dense non-rigid structure-from-motion: A grassmannian perspective. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 254–263, 2018.
- [57] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Monocular dense 3d reconstruction of a complex dynamic scene from two perspective frames. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4649–4657, 2017.
- [58] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Superpixel soup: Monocular dense 3d reconstruction of a complex dynamic scene. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.

- [59] Suryansh Kumar, Ram Srivatsav Ghorakavi, Yuchao Dai, and Hongdong Li. Dense depth estimation of a complex dynamic scene without explicit 3d motion estimation. *arXiv preprint arXiv:1902.03791*, 2019.
- [60] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *International journal of computer vision*, 38(3):199–218, 2000.
- [61] Holger Lange. Advances in the cooperation of shape from shading and stereo vision. In *Second International Conference on 3-D Digital Imaging and Modeling (Cat. No. PR00062)*, pages 46–58. IEEE, 1999.
- [62] Marc Levoy. Display of surfaces from volume data. *IEEE Computer graphics and Applications*, 8(3):29–37, 1988.
- [63] Maxime Lhuillier and Long Quan. A quasi-dense approach to surface reconstruction from uncalibrated images. *IEEE transactions on pattern analysis and machine intelligence*, 27(3):418–433, 2005.
- [64] 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.
- [65] Min Li, Zhenglong Zhou, Zhe Wu, Boxin Shi, Changyu Diao, and Ping Tan. Multi-view photometric stereo: A robust solution and benchmark dataset for spatially varying isotropic materials. *IEEE Transactions on Image Processing*, 29:4159–4173, 2020.
- [66] Zhe Liang, Chao Xu, Jing Hu, Yushi Li, and Zhaopeng Meng. Better together: shading cues and multi-view stereo for reconstruction depth optimization. *IEEE Access*, 8:112348–112356, 2020.
- [67] David B Lindell, Julien NP Martel, and Gordon Wetzstein. AutoInt: Automatic integration for fast neural volume rendering. *arXiv preprint arXiv:2012.01714*, 2020.
- [68] Alex Locher, Michal Perdoch, and Luc Van Gool. Progressive prioritized multi-view stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3244–3252, 2016.
- [69] Fotios Logothetis, Roberto Mecca, and Roberto Cipolla. A differential volumetric approach to multi-view photometric stereo. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1052–1061, 2019.
- [70] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4):163–169, 1987.
- [71] Keyang Luo, Tao Guan, Lili Ju, Haipeng Huang, and Yawei Luo. P-mvsnet: Learning patch-wise matching confidence aggregation for multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10452–10461, 2019.
- [72] Robert Maier, Kihwan Kim, Daniel Cremers, Jan Kautz, and Matthias Nießner. Intrinsic3d: High-quality 3d reconstruction by joint appearance and geometry optimization with spatially-varying lighting. In *Proceedings of the IEEE international conference on computer vision*, pages 3114–3122, 2017.
- [73] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. *arXiv preprint arXiv:2008.02268*, 2020.
- [74] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995.
- [75] Davide Menini, Suryansh Kumar, Martin R Oswald, Erik Sandstrom, Cristian Sminchisescu, and Luc Van Gool. A real-time online learning framework for joint 3d reconstruction and semantic segmentation of indoor scenes. *arXiv preprint arXiv:2108.05246*, 2021.
- [76] 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. *arXiv preprint arXiv:2003.08934*, 2020.
- [77] Daisuke Miyazaki, Kenji Hara, and Katsushi Ikeuchi. Median photometric stereo as applied to the segonko tumulus and museum objects. *International Journal of Computer Vision*, 86(2-3):229, 2010.
- [78] Theo Moons, Luc Van Gool, and Maarten Vergauwen. *3D reconstruction from multiple images: principles*. Now Publishers Inc, 2009.
- [79] MG-H Mostafa, Sameh M Yamany, and Aly A Farag. Integrating shape from shading and range data using neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 15–20. IEEE, 1999.
- [80] Yasuhiro Mukaigawa, Yasunori Ishii, and Takeshi Shikunaga. Analysis of photometric factors based on photometric linearization. *JOSA A*, 24(10):3326–3334, 2007.
- [81] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Chakravarty R Alla Chaitanya, Anton Kaplanyan, and Markus Steinberger. Donerf: Towards real-time rendering of neural radiance fields using depth oracle networks. *arXiv preprint arXiv:2103.03231*, 2021.
- [82] Diego Nehab, Szymon Rusinkiewicz, James Davis, and Ravi Ramamoorthi. Efficiently combining positions and normals for precise 3D geometry. *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH 2005)*, 24(3):536–543, 2005.
- [83] Jaesik Park, Sudipta N Sinha, Yasuyuki Matsushita, Yu-Wing Tai, and In So Kweon. Robust multiview photometric stereo using planar mesh parameterization. *IEEE transactions on pattern analysis and machine intelligence*, 39(8):1591–1604, 2016.
- [84] Hanspeter Pfister. Hardware-accelerated. *Visualization Handbook*, 2(3), 2005.
- [85] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. *arXiv preprint arXiv:2011.13961*, 2020.
- [86] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310. PMLR, 2019.

- [87] Konstantinos Rematas, Ricardo Martin-Brualla, and Vittorio Ferrari. Sharf: Shape-conditioned radiance fields from a single view. In *ICML*, 2021.
- [88] Jieji Ren, Zhenxiong Jian, Xi Wang, Ren Mingjun, Limin Zhu, and Xiangqian Jiang. Complex surface reconstruction based on fusion of surface normals and sparse depth measurement. *IEEE Transactions on Instrumentation and Measurement*, 70:1–13, 2021.
- [89] 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.
- [90] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016.
- [91] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [92] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision*, pages 501–518. Springer, 2016.
- [93] Steven M Seitz and Charles R Dyer. Photorealistic scene reconstruction by voxel coloring. *International Journal of Computer Vision*, 35(2):151–173, 1999.
- [94] B. Shi, Z. Mo, Z. Wu, D. Duan, S. Yeung, and P. Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):271–284, 2019.
- [95] Boxin Shi, Ping Tan, Yasuyuki Matsushita, and Katsushi Ikeuchi. Bi-polynomial modeling of low-frequency reflectances. *IEEE transactions on pattern analysis and machine intelligence*, 36(6):1078–1091, 2013.
- [96] Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung, and Ping Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3707–3716, 2016.
- [97] William M Silver. *Determining shape and reflectance using multiple images*. PhD thesis, Massachusetts Institute of Technology, 1980.
- [98] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. *arXiv preprint arXiv:2012.03927*, 2020.
- [99] Christoph Strecha, Rik Fransens, and Luc Van Gool. Combined depth and outlier estimation in multi-view stereo. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2394–2401. IEEE, 2006.
- [100] Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.
- [101] Tatsunori Taniai and Takanori Maehara. Neural inverse rendering for general reflectance photometric stereo. In *International Conference on Machine Learning (ICML)*, pages 4857–4866, 2018.
- [102] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment—a modern synthesis. In *International workshop on vision algorithms*, pages 298–372. Springer, 1999.
- [103] George Vogiatzis, Philip HS Torr, and Roberto Cipolla. Multi-view stereo via volumetric graph-cuts. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 391–398. IEEE, 2005.
- [104] Eric A Wan and Alex T Nelson. Dual extended kalman filter methods. *Kalman filtering and neural networks*, 123, 2001.
- [105] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patchmatchnet: Learned multi-view patchmatch stereo. *arXiv preprint arXiv:2012.01411*, 2020.
- [106] Robert J Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):191139, 1980.
- [107] Changchang Wu et al. Visualsfm: A visual structure from motion system. 2011.
- [108] Lun Wu, Arvind Ganesh, Boxin Shi, Yasuyuki Matsushita, Yongtian Wang, and Yi Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *Asian Conference on Computer Vision*, pages 703–717. Springer, 2010.
- [109] Tai-Pang Wu and Chi-Keung Tang. Photometric stereo via expectation maximization. *IEEE transactions on pattern analysis and machine intelligence*, 32(3):546–560, 2009.
- [110] Qingshan Xu and Wenbing Tao. Multi-scale geometric consistency guided multi-view stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5483–5492, 2019.
- [111] Qingshan Xu and Wenbing Tao. Learning inverse depth regression for multi-view stereo with correlation cost volume. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12508–12515, 2020.
- [112] Qingshan Xu and Wenbing Tao. Pvsnet: Pixelwise visibility-aware multi-view stereo network. *arXiv preprint arXiv:2007.07714*, 2020.
- [113] Youze Xue, Jiansheng Chen, Weitao Wan, Yiqing Huang, Cheng Yu, Tianpeng Li, and Jiayu Bao. Mvscrf: Learning multi-view stereo with conditional random fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4312–4321, 2019.
- [114] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 767–783, 2018.
- [115] Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, and Long Quan. Recurrent mvsnet for high-resolution multi-view stereo depth inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5525–5534, 2019.
- [116] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Ronen Basri, and Yaron Lipman. Multiview neu-

- ral surface reconstruction with implicit lighting and material. *arXiv preprint arXiv:2003.09852*, 2020.
- [117] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. *arXiv preprint arXiv:2103.14024*, 2021.
- [118] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. *arXiv preprint arXiv:2012.02190*, 2020.
- [119] Lap-Fai Yu, Sai-Kit Yeung, Yu-Wing Tai, and Stephen Lin. Shading-based shape refinement of rgb-d images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1415–1422, 2013.
- [120] Zehao Yu and Shenghua Gao. Fast-mvsnet: Sparse-to-dense multi-view stereo with learned propagation and gauss-newton refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1949–1958, 2020.
- [121] Jure Zbontar, Yann LeCun, et al. Stereo matching by training a convolutional neural network to compare image patches. *J. Mach. Learn. Res.*, 17(1):2287–2318, 2016.
- [122] Li Zhang, Brian Curless, and Steven M Seitz. Rapid shape acquisition using color structured light and multi-pass dynamic programming. In *Proceedings. First International Symposium on 3D Data Processing Visualization and Transmission*, pages 24–36. IEEE, 2002.
- [123] Li Zhang, Brian Curless, and Steven M Seitz. Spacetime stereo: Shape recovery for dynamic scenes. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages II–367. IEEE, 2003.
- [124] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [125] Runze Zhang, Siyu Zhu, Tian Fang, and Long Quan. Distributed very large scale bundle adjustment by global camera consensus. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [126] 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.
- [127] Siyu Zhu, Runze Zhang, Lei Zhou, Tianwei Shen, Tian Fang, Ping Tan, and Long Quan. Very large-scale global sfm by distributed motion averaging. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4568–4577, 2018.
- [128] Michael Zollhöfer, Angela Dai, Matthias Innmann, Chenglei Wu, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Shading-based refinement on volumetric signed distance functions. *ACM Transactions on Graphics (TOG)*, 34(4):1–14, 2015.
- [129] Michael Zollhöfer, Patrick Stotko, Andreas Görlitz, Christian Theobalt, Matthias Nießner, Reinhard Klein, and Andreas Kolb. State of the art on 3d reconstruction with rgb-d cameras. In *Computer graphics forum*, volume 37, pages 625–652. Wiley Online Library, 2018.