

Uncertainty-Aware Deep Multi-View Photometric Stereo

Berk Kaya¹ Suryansh Kumar^{1*} Carlos Oliveira¹ Vittorio Ferrari² Luc Van Gool^{1,3}
ETH Zürich¹, Google Research², KU Leuven³

Abstract

This paper presents a simple and effective solution to the longstanding classical multi-view photometric stereo (MVPS) problem. It is well-known that photometric stereo (PS) is excellent at recovering high-frequency surface details, whereas multi-view stereo (MVS) can help remove the low-frequency distortion due to PS and retain the global geometry of the shape. This paper proposes an approach that can effectively utilize such complementary strengths of PS and MVS. Our key idea is to combine them suitably while considering the per-pixel uncertainty of their estimates. To this end, we estimate per-pixel surface normals and depth using an uncertainty-aware deep-PS network and deep-MVS network, respectively. Uncertainty modeling helps select reliable surface normal and depth estimates at each pixel which then act as a true representative of the dense surface geometry. At each pixel, our approach either selects or discards deep-PS and deep-MVS network prediction depending on the prediction uncertainty measure. For dense, detailed, and precise inference of the object's surface profile, we propose to learn the implicit neural shape representation via a multilayer perceptron (MLP). Our approach encourages the MLP to converge to a natural zero-level set surface using the confident prediction from deep-PS and deep-MVS networks, providing superior dense surface reconstruction. Extensive experiments on the DiLiGenT-MV benchmark dataset show that our method provides high-quality shape recovery with a much lower memory footprint while outperforming almost all of the existing approaches.

1. Introduction

In the coming decade, dense 3D data acquisition of objects is likely to become one of the most important problems in computer vision and industrial machine vision. Moreover, it can be helpful for a wide range of other cutting-edge scientific disciplines such as metrology [12], geometry processing [4], forensics [44], etc. At present, it is widely accepted that methods such as structure-from-motion [21, 55],

multi-view stereo [12], photometric stereo [30, 54, 61], and other standalone approaches [34, 35, 43, 47, 60] are not sufficient on their own to provide detailed and precise 3D reconstruction for all kinds of surfaces [44]. Therefore, methods that combine complementary surface estimates by leveraging more than one modality are often preferred [36, 46].

Among the passive 3D shape acquisition methods, multi-view stereo (MVS) has become the most popular approach [13, 55, 62], especially after the proliferation of cheap digital cameras for high-quality imaging. Yet, MVS works best for Lambertian textured surfaces and gives unreliable results for non-textured objects with non-Lambertian surface reflectance property. Moreover, high-frequency surface details such as indentations and scratches are difficult to recover using MVS methods (see Fig. 1(a)).

On the other hand, photometric stereo (PS) is magnificent at recovering high-frequency surface details using light-varying images [61]. It is also effective for non-textured, and non-Lambertian surfaces [6]. PS allows the recovery of per-pixel depth of the object by integration of the estimated surface normals [24]. However, PS suffers from the main shortcoming: The recovered surface profile is globally deformed by a low-frequency distortion [46]. Such distortion is likely due to numerical integration of the surface normal map without explicit constraints between multiple disconnected regions of the object's surface [46, 52, 64] or non-directional lighting effects (see Fig. 1(b)).

Combining the complementary responses of MVS and PS is known as multi-view photometric stereo (MVPS) [22]. In this paper, we propose an approach that can effectively exploit such complementary surface information. Our work leverages recent advances in deep neural networks. To this end, we use a PatchMatch-based deep-MVS network [59] to infer the per-pixel depth and a CNN-based deep-PS network [27] to infer per-pixel surface normal. But, we know that such deep network models have their accuracy limits and can predict erroneous depth or surface normals in certain parts of the object. In that case, if we naively combine the output predicted by these networks, we may end up having a bad overall result. To resolve this, we extend the deep-PS and deep-MVS networks with per-pixel uncertainty estimation capability. Using the prediction uncertainty as a measure, we select and combine only reliable

*Corresponding Author (k.sur46@gmail.com)

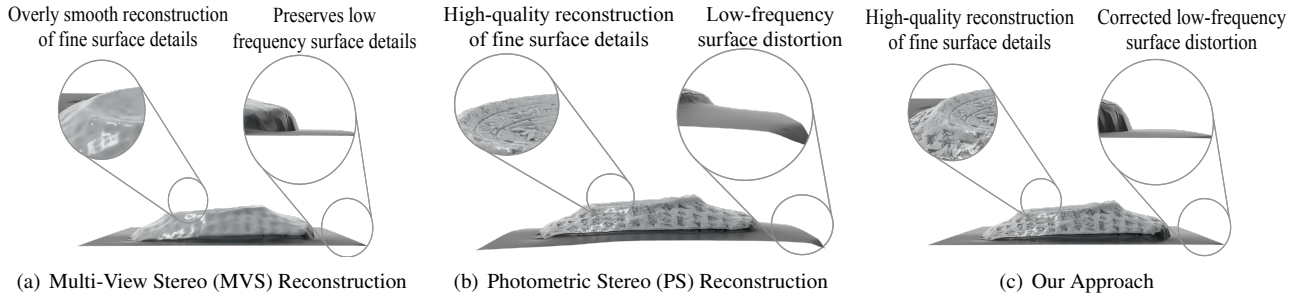


Figure 1. Our method handles the high and low-frequency surface components quite well. It overcomes the high and low-frequency surface reconstruction problem by suitably utilizing the complementary surface estimates via uncertainty modeling, and neural level set optimization (a) MVS reconstruction preserves the plane geometry but loses the finer details. (b) PS captures the fine geometric details but introduces global distortion. (c) Our approach.

surface estimates at each pixel, that is either deep-PS normal, deep-MVS depth, both, or none of the prediction.

Using our approach of selecting and discarding surface estimates may result in the loss of some pixels’ corresponding 3D surface details. To recover those lost geometric details, we introduce a neural network (MLP layers only) based optimization to recover the overall dense shape from those selected surface predictions by representing the object’s shape as level sets of the neural network. Our overall loss function optimization encourages the zero-level set of the neural network to converge to the confident surface estimates. To that end, we first convert the depth estimate to point cloud while keeping the predicted surface normals representation as it is. Our approach then optimizes for parameters of an MLP so that it approximates a signed-distance-function (SDF) to a plausible surface based on the point cloud, surface normals, and an implicit geometric regularization term developed on the Eikonal partial differential equation [8]. Fig.1(c) shows an example reconstruction using our approach. In summary, we make the following contributions:

- We present an effective and easy-to-use deep neural network-based solution to the classical MVPS problem for dense, detailed, and precise recovery of 3D shapes.
- We introduce uncertainty-aware deep-PS and deep-MVS modeling in the MVPS pipeline. Modeling uncertainty helps as a measure in automatic discarding of unreliable surface estimates at a pixel, hence improving robustness.
- We propose an implicit neural shape representation based on the Eikonal term in the MLP loss for natural zero level set surface recovery [8]. It reliably infers the object’s dense surface geometry defined by the confident deep-PS and deep-MVS prediction with better memory foot-print than the methods based on mesh processing [36,49].

2. Related Work

To better place the previous works, we have divided the important MVPS methods literature into two groups.

(a) **Traditional MVPS Methods.** MVPS is a classical 3D

shape acquisition setup introduced by Hernandez et al. [22]. The proposed setup is composed of a turn-table with an object placed at the center of the table for multi-view and photometric stereo image acquisition. The MVPS algorithm proposed by Hernandez et al. [22] combines the multi-view PS results and corrects its low-frequency surface distortion via multi-view geometric constraint. Yet, the method works well only for specific parametric BRDF models [36]. Later, Park et al. [49,50] proposed an uncalibrated MVPS method to recover fine geometric details of the shape. It requires an initial coarse mesh with a 2D displacement map for optimization leading to shape recovery. Still, the method cannot handle objects with diverse surface reflectance properties. Further, it often fails on textureless regions with non-Lambertian reflectances [36]. Logothetis et al. [38] used a volumetric approach to solve MVPS via a variational framework. Recently, Li et al. [36] proposed a systematic geometric approach to MVPS showing state-of-the-art (SOTA) results. However, it consists of several carefully crafted explicit geometric modeling steps such as iso-depth contour estimation, tracing contours, multi-view depth propagation, point sorting, shape optimization, etc. It requires a successful execution of each of these steps applied in a sequel to recover the surface. Hence, re-implementing such an approach is complex, strained, and time-consuming. Furthermore, the use of classical PS and MVS in their pipeline has its limitations; for e.g., classical PS, MVS may not handle a wide range of objects with non-Lambertian properties, etc.

On the contrary, we propose a novel, simple, and everyday deep neural network based solution to MVPS that can handle objects with different reflectance properties and delivers 3D reconstruction accuracy as good as the complex SOTA when tested on the benchmark dataset [36].

(b) **Deep Learning based MVPS Methods.** Kaya et al. [31] recently proposed a neural radiance fields-based approach to solve MVPS. It uses a pre-trained deep-PS model to predict the surface normal. It then conditions the multi-view image rendering on the predicted PS surface normal to recover the surface geometry. Even though the method is

simple and usable, the recovered geometry is poor in quality. Our literature review shows a lack of a robust modern neural network approach to solve MVPS, which is excellent at learning the object’s surface properties from data. Thus, it has become increasingly evident that a simple and effective learning method is essential for the MVPS problem.

Other related work uses an active 3D sensing modality with PS. For instance, Nehab et al. [46] used a structured lighted scanner, whereas Chatterjee et al. [5] relied on a RGB-D sensor to measure the 3D position data. Instead, our work focuses on the classical MVPS setup [22, 36, 50]. It has some apparent advantages over active 3D scanning methods. Firstly, it is easy and cost-effective to perform high-quality image acquisition, as regular cameras are sufficient. Secondly, it is relatively noise-free and gives dense per-pixel information compared to incomplete range data with outliers provided by structured light [18], 3D laser scanner [11], and depth sensors [72].

3. Proposed Approach

We denote $\mathcal{I}^v = \{I_1^v, \dots, I_{N_p}^v\}$ as the set of N_p input PS images for a given view¹ $v \in \{1, \dots, N_m\}$, where N_m is the number of views. For MVS, we follow Li et al. [36] work, which take the median of all PS images per camera view to have MVS images. Concretely, $Y^v = \text{median}(\mathcal{I}^v)$ gives us the set $\mathcal{Y} = \{Y^1, \dots, Y^v, \dots, Y^{N_m}\}$ of MVS images. We assume a calibrated setting, i.e., light source directions, intensities, and camera calibrations are known.

The rest of the section is organized as follows: In Sec. §3.1, we introduce our uncertainty-aware deep-PS network. Next, in Sec. §3.2, we describe the uncertainty-aware deep-MVS network pipeline. Finally, in Sec. §3.3, we explain our neural shape representation approach and level set optimization for dense, detailed, and accurate 3d surface recovery from high-fidelity surface normals and depth estimates.

3.1. Uncertainty-Aware Deep Photometric Stereo

When an object surface is illuminated by k^{th} point light source located in the direction $l_k \in \mathbb{R}^{3 \times 1}$, then the PS image $I_k^v \in \mathbb{R}^{p \times 1}$ (vectorized form) captured by a camera in the view direction $\mathbf{v} \in \mathbb{R}^{3 \times 1}$ w.r.t the object is modeled as:

$$I_k^v = e_k \cdot \Psi(\mathbf{N}_v, l_k, \mathbf{v}) \cdot \max(\mathbf{N}_v^T l_k, 0) + \epsilon_k. \quad (1)$$

Here, p symbolises the total number of object surface pixels, $e_k \in \mathbb{R}_+$ denotes the light intensity, ϵ_k is an additive error. The image formation model in Eq. (1) assumes a real-world object with a reflective surface, whose appearance is encoded by a BRDF Ψ with surface normal $\mathbf{N}_v \in \mathbb{R}^{3 \times p}$ and $\max(\mathbf{N}_v^T l_k, 0)$ accounts for the attached shadows. The formulation in Eq.(1) has led to outstanding developments in PS for recovering fine surface details [57, 61, 63]. Yet,

¹In a turn-table setup, each view captures different part of the object.

modeling unknown reflectance of different objects remains a challenge. Recently, deep learning-based methods have been proposed to utilize neural networks’ ability to learn complicated BRDFs from input data [6, 27, 30, 53, 56, 58]. Accordingly, we adhere to using a supervised deep learning framework to have high-fidelity surface normal predictions at test time for a diverse set of objects with different material properties. To that end, we adopt an observation map based modeling in deep-PS [27] due to its simplicity and notable performance on PS benchmark datasets [56, 57].

• **Observation Map.** An observation map is a 2D matrix that stores normalized grayscale pixel values observed under different light sources (Fig.2 bottom left). We independently create observation map for every object pixel. In a PS setup, there exists a one-to-one mapping between k^{th} light source direction $[l_k(x), l_k(y), l_k(z)]^T \in \mathbb{R}^{3 \times 1}$ and its corresponding $x - y$ coordinate projection $[l_k(x), l_k(y)]^T \in \mathbb{R}^{2 \times 1}$. Utilizing these notations, observation map $\Omega_i^v \in \mathbb{R}^{\omega \times \omega}$ for i^{th} pixel captured from view v is defined as:

$$\Omega_i^v \left(\zeta \left(\omega \cdot \frac{(l_k(x) + 1)}{2} \right), \zeta \left(\omega \cdot \frac{(l_k(y) + 1)}{2} \right) \right) = \frac{I_k^v(i)}{\eta_i e_k} \quad (2)$$

where, ω is the 2D resolution of the observation map and the scalar $\eta_i = \max(I_1^v(i)/e_1, \dots, I_{N_p}^v(i)/e_{N_p})$ is the normalizing constant. Here, $\zeta : \mathbb{R} \mapsto \mathbb{Z}_0^+$, since the projected source direction vectors can take values from $[-1, 1]$, they are scaled to suitable positive integer (including zero).

• **Deep-PS Network Architecture.** The deep-PS network takes Ω_i^v of each pixel as input and regresses to corresponding ground-truth surface normal at train time to obtain deep-PS model (see Fig.2). The network design comprises an initial 3×3 convolution layer, which converts the observation map into a feature block. Next, the architecture has two dense blocks with a transition layer in between. The dense block is composed of a ReLU and 3×3 convolution layer. The transition layer comprises a ReLU, 1×1 convolution, and an average pooling layer. After the second dense block, one convolution and two fully connected layers are applied, followed by a normalization operation to recover unit length surface normals as output. For more details refer [27].

• **Uncertainty Modeling in Deep-PS Network.** The above deep-PS network is not apt for estimating the uncertainty of the predicted surface normals at test time. For our problem, it is imperative to have that information as perfect prediction is not always possible. Accordingly, we modify the deep-PS architecture to provide uncertainty of the predicted surface normals by leveraging the Bayesian neural network (NN) approach [41, 45]. Generally, Bayesian NNs are a simple extension of NNs by placing a prior distribution (generally Gaussian) over the NN’s weights. Compared to standard NN, it has the advantage of providing an uncertainty measure of the network prediction [17, 51]. For completeness, let’s briefly review the standard Bayesian NN framework.

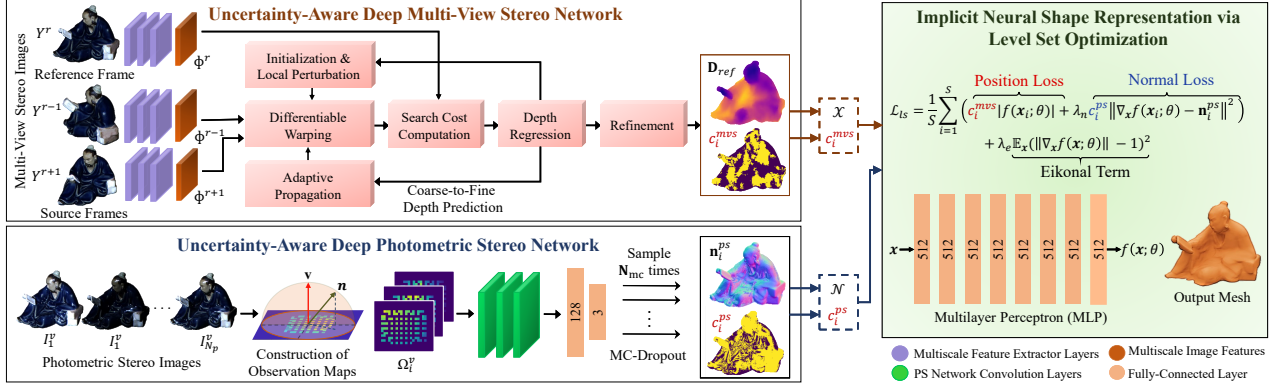


Figure 2. An overview of our uncertainty-aware deep-MVPS approach. We first predict per-pixel depth and surface normals via deep-MVS and deep-PS networks. Then, we recover dense, detailed 3D shape by using these surface estimates in neural level set optimization. Our approach utilizes prediction uncertainty as a measure of reliability and uses highly confident surface estimates in optimization for better surface recovery (Best View on Screen).

Let $\{\mathbf{A}, \mathbf{B}\}$ be the training dataset with \mathbf{A}, \mathbf{B} as input and output sets, respectively. Assume, a Bayesian NN with L layers parameterized by weight $\mathbf{w} = \{\mathbf{W}_j\}_{j=1}^L$ with \mathbf{W}_j as the weight matrix for layer j . The predictive distribution $P(\mathbf{b}^* | \mathbf{a}^*, \mathbf{A}, \mathbf{B})$ for a new input \mathbf{a}^* is formulated as $P(\mathbf{b}^* | \mathbf{a}^*, \mathbf{A}, \mathbf{B}) = \int P(\mathbf{b}^* | \mathbf{a}^*, \mathbf{w}) P(\mathbf{w} | \mathbf{A}, \mathbf{B}) d\mathbf{w}$. However, to determine predictive distribution value is intractable as $P(\mathbf{w} | \mathbf{A}, \mathbf{B})$ is hard to solve analytically [15]. Generally, variational inference (VI) is used to approximate $P(\mathbf{w} | \mathbf{A}, \mathbf{B})$. By introducing the variational distribution $Q_\gamma(\mathbf{w})$ parameterized by γ , the KL divergence between $Q_\gamma(\mathbf{w})$ and $P(\mathbf{w} | \mathbf{A}, \mathbf{B})$ is minimized as $\mathcal{L}_{VI} = -\int Q_\gamma(\mathbf{w}) \log(\mathbf{B} | \mathbf{A}, \mathbf{w}) d\mathbf{w} + KL(Q_\gamma(\mathbf{w}) || P(\mathbf{w}))$.

While VI based on KL divergence is widely used, the assumption of Gaussian distribution on the network parameters increases the complexity of the model. That, in turn, reduces the model efficiency without a significant gain in predictive power. Not long ago, Gal et al. [14] proposed a Bernoulli distribution-based VI approach which provides a simpler and fruitful approximation to the posterior distribution. To adopt Bernoulli distribution approach to our deep-PS network, we model $\mathbf{W}_j = \mathbf{M}_j \cdot \text{diag}([z_{j,u}]_{u=1}^{K_j})$ and $z_{j,u}$ as Bernoulli(q_j) for $j = 1, \dots, L$, $u = 1, \dots, K_{j-1}$. Here, \mathbf{M}_j denotes variational parameters of $\mathbf{W}_j \in \mathbb{R}^{K_j \times K_{j-1}}$ with K_j the number of units at layer j , and $z_{j,u}$ denotes the Bernoulli random variables with probability q_j .

With such parameterization, the integral term in \mathcal{L}_{VI} is approximated by sampling \mathbf{W} from a Bernoulli distribution, also known as Monte Carlo (MC) integration approach. Likewise, the KL divergence term in \mathcal{L}_{VI} is replaced with a weight decay on network parameters [14]. Using these approximations, we train our uncertainty-aware deep-PS network using the following loss function:

$$\mathcal{L}_{ps} = \frac{1}{N_{mc}} \sum_{j=1}^{N_{mc}} \|\tilde{\mathbf{n}}_j - \mathbf{n}_{gt}\|_2^2 + \lambda_w \sum_{j=1}^L \|\mathbf{W}_j\|_2^2 \quad (3)$$

where, N_{mc} is the number of MC samples, $\tilde{\mathbf{n}}_j$ is the estimated surface normal in each forward pass and \mathbf{n}_{gt} is the ground-truth surface normal. As shown in Gal et al. [14] work, we can realize the approximation to Bernoulli distribution by introducing dropout layers in the neural network. Accordingly, we apply dropout with $q_j = 0.2$ after each convolution and fully-connected layer in deep-PS network. We keep dropout layers active during train and test time.

Due to the introduction of dropout layers, we now have a stochastic network. At test time, we run the trained model multiple times, recording the (potentially varying) surface normal prediction at each i^{th} pixel. We then calculate the mean and variance of these multiple predictions at every pixel. The mean is taken as the final prediction $\mathbf{n}_i^{ps} \in \mathbb{R}^{3 \times 1}$ and the variance as its uncertainty $\sigma_i^2 \in \mathbb{R}^{3 \times 1}$. Since our approach is focused on highly confident predictions, we convert the per-pixel variance to a single binary variable c_i^{ps} . We set $c_i^{ps} = 1$, if $\|\sigma_i^2\|_1 < \tau_{ps}$ and $c_i^{ps} = 0$, otherwise.

3.2. Uncertainty-Aware Deep Multi-View Stereo

Similar to deep-PS network, we aim to have a per-pixel uncertainty measure but now on the depth prediction. One natural way is to similarly use MC dropout strategy to deep-MVS network. Fortunately, there already exist deep MVS frameworks which have the intrinsic ability to implicitly provide uncertainty measure via confidence values of their depth predictions [7, 25, 59, 69]. Hence, it is inefficient to add extra complexity by introducing MC dropout layers. Among [7, 25, 59, 69], we use [59], i.e., PatchMatch based deep-MVS network due to its recent SOTA performance on MVS benchmarks and fast inference on large scale images.

• **PatchMatch based Deep-MVS Network.** Similar to PatchMatch algorithm [2], the PatchMatch based deep-MVS network employs [2] via similar three steps (but in 3d scene space) as follows: (i) Initialization step: Generating depth hypotheses, (ii) Propagation step: Propagate the hypotheses to neighbors, and (iii) Evaluation step: Compute

the similarity cost and search for best solution. We apply these steps on per-pixel multi-scale features that are hierarchically extracted from \mathcal{Y} at M different resolution scales [37, 59] (see Fig.2 top left). This allowed us to estimate depth in a coarse-to-fine manner. Before we introduce the steps, let’s denote the reference frame by $Y^r \in \mathbb{R}^{w \times h}$, coordinates of the i^{th} pixel by \mathbf{y}_i , frame r feature by Φ^r , and camera r intrinsic calibration matrix by \mathbf{K}_r . For each reference frame, we pick N_s source frames where $Y^s \in \mathbb{R}^{w \times h}$ denotes a source frame. $(\mathbf{R}_{r,s}, \mathbf{t}_{r,s})$ denotes the relative motion between frame r and s . We skip to add extra notation for stage number for simplicity of writing.

(i) Initialization. In the first iteration, per-pixel \mathcal{D}_f depth hypotheses are sampled. Once initialized, “local perturbations” are invoked in the subsequent iteration at each stage to diversify the hypotheses [3]. For local perturbations, per-pixel N_l^m hypotheses are generated at stage m .

(ii) Propagation. Existing hypotheses are enriched using spatially neighboring pixels that are likely to have similar depth values. For that, an “adaptive propagation” approach that uses a learnable offset to gather hypotheses from the same physical surface rather than a fixed set of neighbors is used. Accordingly, 2D CNN is applied on Φ^r to learn a 2D offset for each pixel and obtain N_d^m additional hypotheses at stage m using the depth map of the previous iteration.

(iii) Evaluation. The best solution is searched by evaluating the similarity cost for existing depth hypotheses. For that, group-wise correlation between $\Phi^r(\mathbf{y}_i)$ and warped source feature $\Phi^s(\mathbf{y}_i^{s,j})$ is calculated for every pixel and depth hypothesis [67]. Here, group-wise correlation is obtained by dividing the features into groups along channel dimension and computing the inner product per group. $\mathbf{y}_i^{s,j}$ is the warped coordinates of the i^{th} pixel at source view s and is computed with “differentiable warping” relation $\mathbf{y}_i^{s,j} = \mathbf{K}_s(\mathbf{R}_{r,s}(d_j(\mathbf{y}_i) \cdot \mathbf{K}_r^{-1}\mathbf{y}_i) + \mathbf{t}_{r,s})$ where $d_j(\mathbf{y}_i)$ stands for the j^{th} depth hypothesis at pixel coordinates \mathbf{y}_i .

The group-wise correlation costs over number of views are aggregated with per-pixel view weight [55, 68]. Then, a 3D convolution layer with $1 \times 1 \times 1$ kernel is applied on the aggregated cost to obtain “search cost” per-pixel and depth hypothesis $\mathbf{J} \in \mathbb{R}^{w \times h \times \mathcal{D}}$. The search cost is further aggregated over a spatial window into $\tilde{\mathbf{J}}(\mathbf{y}_i, j)$ for increased robustness. For per-pixel “depth regression”, softmax (σ) is applied to $\tilde{\mathbf{J}}(\mathbf{y}_i, j)$ and expectation over depth hypotheses is evaluated as $\mathbf{D}(\mathbf{y}_i) = \sum_{j=0}^{\mathcal{D}-1} d_j(\mathbf{y}_i) \cdot \sigma(\tilde{\mathbf{J}}(\mathbf{y}_i, j))$.

Subsequently, per-pixel confidence measure ρ_i is computed using the predicted probability of the most likely depth hypothesis, i.e. $\rho_i = \sigma(\tilde{\mathbf{J}}(\mathbf{y}_i, j^*))$. Finally, an independent depth residual network based on [26] is used for “refinement” to obtain output depth map \mathbf{D}_{ref} . For more details on PatchMatch based deep-MVS network, refer [59].

• MVS Loss Function. We use l_1 loss between the estimated depth and the ground-truth depth at the same resolu-

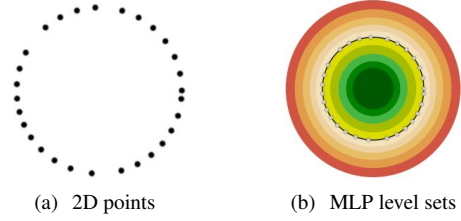


Figure 3. (a) Sparse 2D point cloud. (b) The Eikonal term based implicit geometric regularization in the optimization gives plausible zero level set for (a)—shown in black. The different colors in (b) show the level sets.

tion. The total MVS loss takes into account the PatchMatch loss \mathcal{L}_{pm} at each stage along with refined depth loss \mathcal{L}_{ref} .

$$\mathcal{L}_{mvs} = \mathcal{L}_{pm} + \mathcal{L}_{ref}, \text{ where } \mathcal{L}_{pm} = \sum_{m=1}^M \sum_{t=1}^{N_{iter}^m} \mathcal{L}_t^m \quad (4)$$

Here, N_{iter}^m denotes the total number of iterations at stage m .

• Uncertainty Modeling in Deep-MVS Network. To have the notion of per-pixel depth uncertainty, we convert depth prediction confidence value (ρ_i ’s) to a binary variable c_i^{mvs} . We set $c_i^{mvs} = 1$ if $\rho_i > \tau_{mvs}$ and $c_i^{mvs} = 0$, otherwise.

3.3. Implicit Neural Shape Representation

Once we have the complementary information of the object shape, i.e., surface normals, depth, and the fidelity of prediction at hand, our goal is to combine them effectively for dense surface reconstruction. It is pretty natural to go for regular volumetric fusion (VF) approaches [9, 10, 29, 42] since surface normals can also provide depth by simple integration [23]. However, we know VF uses a fixed size cubic grid independent of the object’s geometry, and therefore, may not obey the geometry of the shapes we want to model. Consequently, we propose to work directly on the confident raw surface estimates. For that, we convert deep-MVS network per-pixel depth prediction to point cloud $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^S \subset \mathbb{R}^3$, where S is the total number object pixel across all views. Additionally, we use per-pixel surface normal prediction from deep-PS network $\mathcal{N} = \{\mathbf{n}_i^{ps}\}_{i=1}^S \subset \mathbb{R}^3$. We propose to recover the 3d surface by optimizing the parameters of a MLP $f(\mathbf{x}; \theta)$. The suggested MLP approximates the signed distance function (SDF) defined by \mathcal{X} and \mathcal{N} . We consider the following loss function for MLP optimization:

$$\mathcal{L}_{ls} = \frac{1}{S} \sum_{i=1}^S \left(\underbrace{c_i^{mvs} |f(\mathbf{x}_i; \theta)|}_{\text{position loss}} + \lambda_n \underbrace{c_i^{ps} \|\nabla_{\mathbf{x}} f(\mathbf{x}_i; \theta) - \mathbf{n}_i^{ps}\|}_{\text{normal loss}} \right) + \lambda_e \underbrace{\mathbb{E}_{\mathbf{x}} (\|\nabla_{\mathbf{x}} f(\mathbf{x}; \theta)\| - 1)^2}_{\text{Eikonal term}} \quad (5)$$

The first term encourages the zero level set to converge to high-fidelity position estimates (i.e., for $c_i^{mvs} = 1$). The

second term forces the local gradients to be consistent with the reliable normal estimates (i.e., for $c_i^{ps} = 1$). The final term stands for the Eikonal regularization, and it is computed by taking the expectation \mathbb{E} over probability distribution $\mathbf{x} \sim \mathcal{P}$. It is noted over several experiments that the Eikonal term is impressive at implicitly regularizing the zero level set. For more details, refer [8, 19].

4. Experiment and Results

Train set. We use CyclesPS synthetic dataset [27] to train our *deep-PS network*. It consists of 15 shapes each of which is rendered with diffuse, specular, and metallic BRDF's using 1300 light sources. We use 90% of the data for training and 10% for validation. For training the *deep-MVS network*, we use DTU MVS dataset [1]. It provides images of 80 scenes captured from 49 or 64 views (depending on the subject) with their ground-truth (GT) depth maps. We keep the training and validation splits same as outlined in [28].

Test set. We used DiLiGenT-MV benchmark dataset [36] as the test set to perform all our experiments, statistical evaluations, and ablations. It comprises real-world objects with complex surface profiles and BRDF properties, making it an ideal choice for MVPS algorithm evaluation. It contains MVS and PS images of five real objects captured from 20 viewpoints using the classical turntable MVPS acquisition setup [22]. For each view, 96 images are acquired, with each image illuminated by a distinct light source. The distance between the object and the camera center is set to $\sim 1.5m$. In addition, the dataset supplies the light source and camera calibration information.

• **Implementation Details.** We implemented our approach in Python 3.8 using PyTorch 1.7.1 library. We conducted all the experiments on a commodity desktop supported with NVIDIA GPU with 11GB of RAM. We trained deep-PS and deep MVS networks independently in a supervised setting.

(i) **Uncertainty-Aware Deep Photometric Stereo.** We first generate pixel-wise observation maps $\Omega_i^v \in \mathbb{R}^{32 \times 32}$ using CyclesPS images to train our deep-PS network. For each observation map, we randomly pick between 50 to 1300 light sources. We train the network for 10 epochs using Adam optimizer [32] with a learning rate of 0.1. During training, we set $N_{mc} = 10$ and $\lambda_w = 10^{-4}$ in our loss function (see Eq.(3)). After training, we perform uncertainty based inference on DiLiGenT-MV images. For that, we first generate observation maps for each pixel. Then, we run the deep-PS network on each observation map 100 times following MC-Dropout approach [14]. We calculate, for each pixel, the mean and variance of the outputs to obtain surface normal \mathbf{n}_i^{ps} and its uncertainty σ_i^2 . We set $\tau_{ps} = 0.03$ to obtain c_i^{ps} ($c_i^{ps} = 1$ if $\|\sigma_i^2\|_1 < \tau_{ps}$ and $c_i^{ps} = 0$, otherwise).

(ii) **Uncertainty-Aware Deep Multi-View Stereo.** The deep-MVS network is trained using DTU MVS dataset [1]. It is trained for 8 epochs using Adam optimizer [32] and learn-

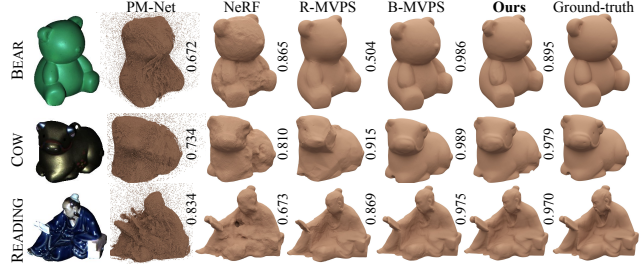


Figure 4. Comparison of the reconstruction quality with PM-Net [59], NeRF [43], R-MVPS [49], and B-MVPS [36] using \mathcal{F} -score metric.

ing rate of 0.001. To predict depth with coarse-to-fine approach, we set number of stages $M = 3$ for performing PatchMatch at different scales with $N_{iter}^3 = 2$, $N_{iter}^2 = 2$, $N_{iter}^1 = 1$ (higher M indicates coarser scale). $\mathcal{D}_f = 48$ depth hypotheses are used at initialization. For local perturbation $N_l^3 = 16$, $N_l^2 = 8$, $N_l^1 = 8$, and $N_d^3 = 16$, $N_d^2 = 8$, $N_d^1 = 0$ at propagation steps. At last, depth refinement is performed at the original image resolution [59].

For testing, we use DiLiGenT-MV images with $N_s = 2$ source images per reference image. We predict per-pixel depth and confidence measure ρ_i for all views using the above parameters. We set $\tau_{mvs} = 0.9$ to obtain c_i^{mvs} , where $c_i^{mvs} = 1$ if $\rho_i > \tau_{mvs}$ and $c_i^{mvs} = 0$ otherwise.

(iii) **Overall Loss Optimization.** We optimize for a zero-level set surface defined by highly confident estimates in \mathcal{X} and \mathcal{N} (see Sec. §3.3). For that, we first perform a precautionary multi-view consistency check to eliminate a few spurious 3d points. MLP with 8 layers is then used on remaining highly-confident estimates to learn suitable implicit neural shape representation. Here, each MLP layer contains 512 hidden units. A skip connection combines the input to the 4th layer to speed up learning [19]. We set the parameters $\lambda_n = 10$ and $\lambda_e = 1$ for our loss function (Eq.(5)). The distribution \mathcal{P} for the expectation in Eq.(5) is taken as average of (a) sum of Gaussian centered at points of \mathcal{X} locally and (b) a uniform distribution globally. The standard deviation of Gaussian at a point is taken as the distance to the 50th nearest neighbor. For optimization, we used Adam optimizer [32] with a learning rate of 0.001. We train the MLP for 10^5 epochs by sampling 2^{14} elements from \mathcal{X}, \mathcal{N} in each batch. We run the trained MLP on a volumetric grid of size 512^3 , which is then used by marching cubes algorithm [39] to extract the mesh corresponding to the zero level set of MLP based neural shape representation.

4.1. Statistical Study

Evaluation Metrics. We evaluated our results using the Chamfer- L_1 , precision, and \mathcal{F} -score metric as defined in [33]. In our evaluation, we used \mathcal{F} -score metric with $1mm$ distance threshold τ . Next, we present our baseline results, which we have divided into three sub-categories.

Method Type →		Deep Multi-View Stereo		View-Synthesis		Photometric Stereo			Ours
Dataset ↓	Method →	MVSNet [69]	PM-Net [59]	NeRF [43]	IDR [70]	Robust PS [48]	SDPS-Net [6]	CNN-PS [27]	
	BEAR	0.135	0.672	0.865	0.053	0.266	0.239	0.293	0.895
	BUDDHA	0.147	0.799	0.713	0.150	0.367	0.298	0.363	0.922
	COW	0.095	0.734	0.810	0.098	0.245	0.447	0.511	0.979
	POT2	0.126	0.666	0.859	0.079	0.231	0.464	0.632	0.907
	READING	0.115	0.834	0.673	0.073	0.242	0.188	0.508	0.970
	AVERAGE	0.124	0.741	0.784	0.091	0.270	0.327	0.461	0.935

Table 1. \mathcal{F} -score comparison with standalone methods on DiLiGenT-MV dataset [36]. The statistics show that our method achieves better results compared to the SOTA standalone multi-view and PS methods, thereby showing the advantage of using the complementary surface information in a MVPS setup.

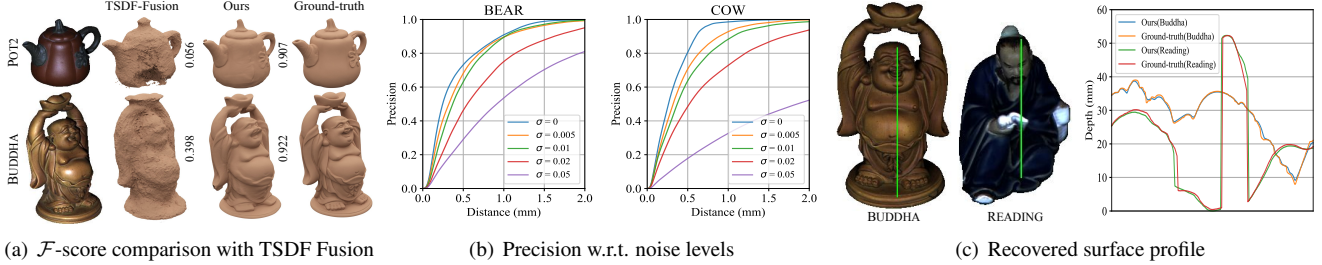


Figure 5. (a) Comparison of the reconstruction quality with TSDF Fusion [9], which is a standard method of choice for robust 3D fusion (outlier removal). (b) Surface reconstruction accuracy for Bear and Cow objects under different noise levels. We report precision as a function of distance threshold (τ) to show the fraction of accurately reconstructed points. (c) Estimated surface profile using our approach, showing how the recovered 3D shape follows the ground-truth surface profile curve when compared across the arbitrarily chosen geodesics.

4.1.1 Baseline Comparisons

(a) **With standalone methods.** Here, we compare our performance with methods that use either PS or MVS setup. Such an experiment helps us understand the benefit of using PS and MVS information together and how accurately we can reconstruct the shape with standalone methods. We used \mathcal{F} -score to compare our method’s performance with SOTA PS, MVS, and view-synthesis methods. Table(1) provides the statistics for the same, indicating the clear advantage of our approach against the standalone methods.

Dataset ↓	Method →	Kaya et al. [31]	R-MVPS [49]	B-MVPS [36]	Ours	Difference with [36]
	BEAR	0.856	0.504	0.986	0.895	0.091
	BUDDHA	0.690	0.935	0.934	0.922	0.012
	COW	0.844	0.915	0.989	0.979	0.010
	POT2	0.858	0.458	0.984	0.907	0.077
	READING	0.720	0.869	0.975	0.970	0.005
	AVERAGE	0.794	0.736	0.974	0.935	0.039

Table 2. \mathcal{F} -score comparison with different MVPS methods on DiLiGenT-MV dataset [36]. The statistics show that our method performs far better than R-MVPS and compares favorably with the SOTA B-MVPS [36]. The point to note is that we can get results close to the SOTA with a simple and easy-to-implement method.

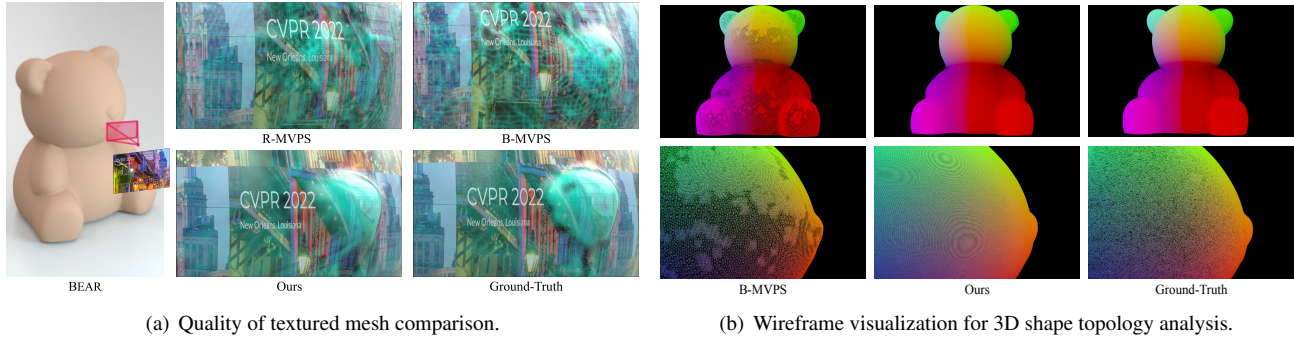
(b) **With MVPS methods.** We compare our method with well-known MVPS methods, i.e., R-MVPS [49], B-MVPS [36], and Kaya et al. [31]. Table(2) provides \mathcal{F} -score comparison with these methods. Our method performs better than R-MVPS [49], Kaya et al. [31] and compares favorably with B-MVPS [36] with the minor difference in \mathcal{F} -score values i.e., less than 10^{-1} in all object categories. Additionally, we want to emphasize again that B-MVPS [36] relies on several carefully crafted explicit geometric steps and refinements that are complex and time-consuming, while our

deep MVPS approach is easy to implement and realize.

(c) **With standard volumetric fusion method.** For estimating the dense 3D surface, it is possible to combine the deep-MVS depth map and depth from PS surface normal map via widely used robust 3D fusion technique, i.e., TSDF fusion [9]. To show that our approach is better at estimating the object’s surface than TSDF fusion, we executed the TSDF fusion algorithm on our deep-MVS depth and depth from deep-PS output. The qualitative and quantitative comparison is presented in Fig.5(a). The results clearly show that TSDF fusion provides undesirable output. On the other hand, our approach takes proper care of the surface estimates and provides much better 3D surface reconstruction.

4.1.2 Ablation Study and Further Analysis

(a) **Effect of uncertainty modeling.** To understand the impact of uncertainty modeling on our method’s performance, we analyzed our results under three distinct settings of overall loss function i.e., Eq.(5): (i) We remove both deep-MVS and deep-PS confidence variable, i.e., c_i^{mvs} , c_i^{ps} from Eq.(5) (w/o uncertainty modeling), (ii) We keep the deep-MVS confidence variable (c_i^{mvs}) and drop the deep-PS confidence variable (c_i^{ps}) from Eq.(5) (w/o PS uncertainty modeling), and (iii) We keep the deep-PS confidence variable (c_i^{ps}) and drop the deep-MVS confidence variable (c_i^{mvs}) from Eq.(5) (w/o MVS uncertainty modeling). Table(3) shows the Chamfer- L_1 metric results obtained under the three settings. The results indicate that incorporating uncertainty information to the loss helps handle the erroneous estimations, and therefore, results in better 3d reconstruction.



(a) Quality of textured mesh comparison.

(b) Wireframe visualization for 3D shape topology analysis.

Figure 6. (a) We transfer the CVPR’22 logo texture on the local region (around the nose) of the mesh recovered using SOTA MVPS methods. It can be observed that the texture pattern on our recovered mesh is closer to GT compared to B-MVPS [36] and R-MVPS [49] (notice the shift of the text). We want to emphasize that if the local topology is same, it must place the text at similar location as can be seen in ours result and GT. (b) Colored Wireframe comparison with SOTA B-MVPS [36]. Clearly, the distribution of geometric primitives on B-MVPS result is uneven and unbalanced compared to ours.

Settings ↓ Dataset →	BEAR	BUDDHA	COW	POT2	READING	AVG
w/o Uncertainty Modeling	0.468	0.485	0.365	0.557	0.380	0.451
w/o PS Uncertainty Modeling	0.443	0.481	0.381	0.484	0.377	0.433
w/o MVS Uncertainty Modeling	0.457	0.473	0.339	0.636	1.024	0.586
Ours (LCNet [6] light)	0.481	0.465	0.346	0.481	0.381	0.431
Ours	0.415	0.455	0.329	0.515	0.355	0.414

Table 3. Effect of uncertainty modeling and light calibration on the reconstruction quality of our approach. The results show Chamfer- L_1 metric (lower is better). The numbers confirm that utilizing the estimated uncertainty of both PS and MVS produces the best results. Further, our approach performs well without the exact light sources i.e., LCNet light sources [6].

(b) Effect of light sources. Here, we study the behavior of our approach under the uncalibrated-PS setting where GT light sources are not given. Instead, we used a pre-trained neural network to have an initial estimate of light sources. Precisely, we used LCNet [6] model to have light source direction and intensity values. Table(3) shows Chamfer- L_1 metric obtained using LCNet predicted light sources information. Our method performs almost equivalently well, showing robustness to small errors in the light calibration.

(c) Effect of noise. We consider the MVPS acquisition setup in our work, where imaging noise is inevitable. So, we analyze the behavior of our approach under imaging noise. To that end, we add zero-mean Gaussian noise to images with different standard deviations (σ 's). Fig.5(b) shows the precision curve of the recovered surface as a function of the distance threshold τ . Precisely, it measures the fraction of points that are reconstructed accurately. The plots show that increasing the noise level degrades the performance. Further, we infer that behavior among subjects varies as signal-to-noise ratio changes. We can observe that our method is robust, and the performance drop is not random.

(d) Quality of reconstructed surface geometry. To perform this experiment, we first analyzed the surface profile of our reconstructed shape across randomly chosen geodesics on the object. Fig.5(c) shows that our recovered surface profile closely follows GT. Next, we performed a local surface analysis of recovered mesh. Although evaluation of recov-

ered shape based on global performance metric is already discussed, it may not reflect the true picture of surface topology since the actual distribution of mesh on GT shape is not known a priori. So, as a second experiment, we transferred texture on a local mesh topology and qualitatively compared the results. Fig.6(a) shows texture transfer results on the mesh recovered using different methods, and clearly, our textured mesh reflects fine text details and appears close to GT. Further, we analyzed the colored Wireframe of the recovered shape compared to SOTA B-MVPS [36]. Fig.6(b) Wireframe model shows that the distribution of geometric primitives in our recovered shape is smooth, regular, and close to GT, whereas B-MVPS [36] has an unbalanced distribution of geometric primitives.

Limitations. This work assumed light sources and camera calibration information are given as input. Further, we tested our method on [36] dataset, which is generally composed of isotropic material objects, and it will be interesting to investigate our method’s performance on anisotropic material objects. Nevertheless, for now, we are limited by the availability of open-source MVPS datasets.

5. Conclusion

This paper explores the field of MVPS using the concepts from deep learning, geometry, and uncertainty. Unlike existing MVPS methods, which treat 3D shape reconstruction as point estimation and geometric optimization problems, we propose learning the fidelity of surface estimation and recovering the shape based on the implicit neural shape representation. Without using complex geometric steps, we observed that our simple neural network based approach could provide results comparable to the best available algorithm. Thus, we believe our work will enable broader use of the MVPS approach in precise 3D data acquisition.

Acknowledgement. This work was supported by Focused Research Award from Google, ETH Zürich Foundation (CVL, ETH 2019-HE-318, 2019-HE-323, 2020-FS-351, 2020-HS-411).

References

- [1] Henrik Aanæs, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders Bjorholm Dahl. Large-scale data for multiple-view stereopsis. *International Journal of Computer Vision*, 120(2):153–168, 2016.
- [2] 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.
- [3] Michael Bleyer, Christoph Rhemann, and Carsten Rother. Patchmatch stereo-stereo matching with slanted support windows. In *Bmvc*, volume 11, pages 1–11, 2011.
- [4] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] Michael G Crandall and Pierre-Louis Lions. Viscosity solutions of hamilton-jacobi equations. *Transactions of the American mathematical society*, 277(1):1–42, 1983.
- [9] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996.
- [10] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5868–5877, 2017.
- [11] 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.
- [12] Yasutaka Furukawa and Carlos Hernández. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015.
- [13] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2009.
- [14] Yarin Gal and Zoubin Ghahramani. Bayesian convolutional neural networks with bernoulli approximate variational inference. *arXiv preprint arXiv:1506.02158*, 2015.
- [15] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.
- [16] 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.
- [17] Jakob Gawlikowski, Cedric Rovic Njietcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *arXiv preprint arXiv:2107.03342*, 2021.
- [18] Jason Geng. Structured-light 3d surface imaging: a tutorial. *Advances in Optics and Photonics*, 3(2):128–160, 2011.
- [19] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of Machine Learning and Systems 2020*, pages 3569–3579. 2020.
- [20] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3273–3282, 2019.
- [21] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [22] Carlos Hernandez, George Vogiatzis, and Roberto Cipolla. Multiview photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3):548–554, 2008.
- [23] B.K. Horn. *Shape Form Shading*. Artificial intelligence. MIT Press, 1989.
- [24] 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.
- [25] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deepmvs: Learning multi-view stereopsis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2821–2830, 2018.
- [26] Tak-Wai Hui, Chen Change Loy, and Xiaoou Tang. Depth map super-resolution by deep multi-scale guidance. In *European conference on computer vision*, pages 353–369. Springer, 2016.
- [27] Satoshi Ikehata. Cnn-ps: Cnn-based photometric stereo for general non-convex surfaces. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [28] Mengqi Ji, Juergen Gall, Haitian Zheng, Yebin Liu, and Lu Fang. Surfacer-net: An end-to-end 3d neural network for multiview stereopsis. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2307–2315, 2017.
- [29] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1251–1261, 2020.
- [30] 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*, pages 3804–3814, 2021.

- [31] Berk Kaya, Suryansh Kumar, Francesco Sarno, Vittorio Ferrari, and Luc Van Gool. Neural radiance fields approach to deep multi-view photometric stereo. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1965–1977, 2022.
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [33] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017.
- [34] 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.
- [35] 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.
- [36] 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.
- [37] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] David JC MacKay. A practical bayesian framework for back-propagation networks. *Neural computation*, 4(3):448–472, 1992.
- [42] Mateusz Michalkiewicz, Jhony K Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Deep level sets: Implicit surface representations for 3d shape inference. *arXiv preprint arXiv:1901.06802*, 2019.
- [43] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020.
- [44] Theo Moons, Luc Van Gool, and Maarten Vergauwen. *3D reconstruction from multiple images: principles*. Now Publishers Inc, 2009.
- [45] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [46] 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.
- [47] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE international symposium on mixed and augmented reality*, pages 127–136. IEEE, 2011.
- [48] Tae-Hyun Oh, Hyeonwoo Kim, Yu-Wing Tai, Jean-Charles Bazin, and In So Kweon. Partial sum minimization of singular values in rpca for low-level vision. In *Proceedings of the IEEE international conference on computer vision*, pages 145–152, 2013.
- [49] 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.
- [50] Jaesik Park, Sudipta N Sinha, Yasuyuki Matsushita, Yu-Wing Tai, and In So Kweon. Multiview photometric stereo using planar mesh parameterization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1161–1168, 2013.
- [51] Nicholas G Polson and Vadim Sokolov. Deep learning: A bayesian perspective. *Bayesian Analysis*, 12(4):1275–1304, 2017.
- [52] Yvain Quéau, Jean-Denis Durou, and Jean-François Aujol. Variational methods for normal integration. *Journal of Mathematical Imaging and Vision*, 60(4):609–632, 2018.
- [53] 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.
- [54] Francesco Sarno, Suryansh Kumar, Berk Kaya, Zhiwu Huang, Vittorio Ferrari, and Luc Van Gool. Neural architecture search for efficient uncalibrated deep photometric stereo. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 361–371, 2022.
- [55] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [56] 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.
- [57] 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.

- [58] Tatsunori Tanai and Takanori Maehara. Neural inverse rendering for general reflectance photometric stereo. In *International Conference on Machine Learning (ICML)*, pages 4857–4866, 2018.
- [59] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patchmatchnet: Learned multi-view patchmatch stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14194–14203, 2021.
- [60] Lawrence B Wolff. Polarization vision: a new sensory approach to image understanding. *Image and Vision computing*, 15(2):81–93, 1997.
- [61] Robert J Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):191139, 1980.
- [62] Changchang Wu et al. Visualsfm: A visual structure from motion system. 2011.
- [63] 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.
- [64] Wuyuan Xie, Miaohui Wang, Mingqiang Wei, Jianmin Jiang, and Jing Qin. Surface reconstruction from normals: A robust dgp-based discontinuity preservation approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5328–5336, 2019.
- [65] Haofei Xu and Juyong Zhang. Aanet: Adaptive aggregation network for efficient stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1959–1968, 2020.
- [66] 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.
- [67] 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.
- [68] Qingshan Xu and Wenbing Tao. Pvsnet: Pixelwise visibility-aware multi-view stereo network. *arXiv preprint arXiv:2007.07714*, 2020.
- [69] 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.
- [70] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33, 2020.
- [71] Enliang Zheng, Enrique Dunn, Vladimir Jovic, and Jan-Michael Frahm. Patchmatch based joint view selection and depthmap estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1510–1517, 2014.
- [72] 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.