

Inferring Super-Resolution Depth from a Moving Light-Source Enhanced RGB-D Sensor: A Variational Approach

Lu Sang Bjoern Haefner Daniel Cremers
 Technical University of Munich, Germany
 {lu.sang,bjoern.haefner,cremers}@tum.de

Abstract

A novel approach towards depth map super-resolution using multi-view uncalibrated photometric stereo is presented. Practically, an LED light source is attached to a commodity RGB-D sensor and is used to capture objects from multiple viewpoints with unknown motion. This non-static camera-to-object setup is described with a nonconvex variational approach such that no calibration on lighting or camera motion is required due to the formulation of an end-to-end joint optimization problem. Solving the proposed variational model results in high resolution depth, reflectance and camera pose estimates, as we show on challenging synthetic and real-world datasets.

1. Introduction

RGB-D sensors are a cheap and easy way to capture RGB and depth images. Yet, the delivered depth is prone to noise, quantization, missing data and a lower resolution compared to its RGB image. The Intel RealSense D415 delivers RGB images almost twice as large in resolution compared to its companion depth images. As the RGB data is substantially better, it seems a good choice to rely on the image intensities to perform depth refinement. Yet, due to its complexity improving a sensors depth map remains an open and challenging problem in the computer vision community. It is often tackled by capturing a sequence of RGB-D frames under multiple view-points. The data, along with the corresponding RGB information can then be used in order to refine the geometric measurements.

A classic technique to estimate geometry with fine scale detail is called Photometric Stereo (PS). It relies on an image sequence where each frame is captured under different illumination. Then, PS estimates fine scale details of the geometry, as well as reflectance and possibly lighting properties of the scene. In classic PS approaches the camera is assumed to be static with respect to the object, while only illumination varies for each frame.

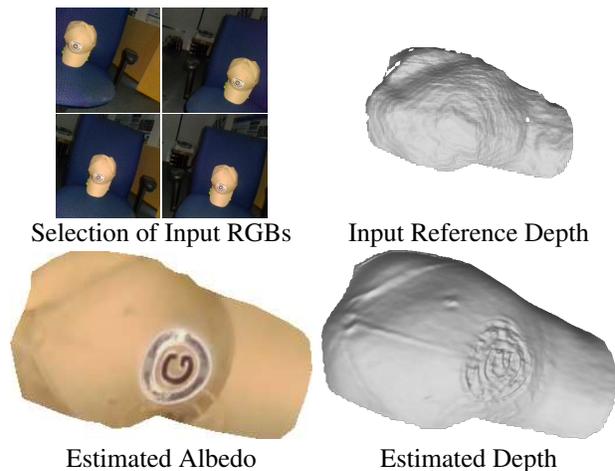


Figure 1: We present a novel method that combines depth super-resolution and multi-view uncalibrated photometric stereo. Our approach is able to recover a refined high-resolution depth and reflectance estimate based on multiple RGB images and its companion depth map acquired with a moving RGB-D camera and an attached LED light source.

For many purposes the classical PS setup is rather impractical, as a moving camera is usually used to capture the object of interest. We propose to close the gap between photometric stereo and camera motion in order to be able to estimate a high resolution depth map, cf. Figure 1. To achieve varying illumination in each frame, we simply attach an LED light source to the RGB-D camera. No further calibration on the light source is needed, as we put forward an end-to-end variational model which estimates camera motion, depth, reflectance, and lighting in an alternating manner.

The rest of the paper is organized as follows: Section 2 reviews related work, while the proposed model is introduced in Section 3. This is followed by the numerical details in Section 4. The model is evaluated in Section 5 on both, synthetic and real-world datasets. Eventually, Section 6 summarizes and suggests future research directions.

2. Related Work

Let us first give some background information on the two problems of *depth super-resolution* and *photometric stereo*.

Depth Super-Resolution. Super-Resolution (SR) is a problem which aims at enhancing the resolution of given imaging data. The two most common data sources are RGB and depth sensors. Thus, one commonly differs between two problem statements: image SR [12, 13, 16, 32, 57, 58] which tries to solve SR for RGB images, and depth SR [31, 34, 37, 48] which considers depth maps as the source of data to be enhanced. The applications of SR range from surveillance [11] and medical imaging [18] to multi-view reconstruction [16] and RGB-D sensing [40, 43]. Our approach is more oriented towards the latter, and we aim to tackle the problem of depth SR which can be largely classified into two categories: single depth map SR and multi depth map SR. The former tries to infer a high resolution depth map from a single observation. One can for instance use an adaption of the non-local means formulation to data typical for depth maps [26]. Another variant is to incorporate the companion RGB image, e.g. to guide a bilateral filtering process and sub-pixel refinement [59] or to use an anisotropic total generalized variation regularization term in the optimization problem [15]. Multi-scale guided convolutional networks can complement low-resolution depth features with high resolution RGB features [27], or combinations of classical variational approaches with deep networks were used [49]. Another variant is to train a network end-to-end to directly learn a mapping between low-resolution and high-resolution depth images [51]. Recently, it was shown that taking into account photometric approaches helps solving the problem of depth super-resolution, due to incorporating the physical relationship between intensity and depth [19, 21, 38, 44]. Our approach builds upon these ideas to recover a fine scale, geometrically detailed depth image based on a methodology called photometric stereo.

Photometric Stereo. Before diving into PS, let us first recall its pioneering work “Shape-from-Shading” [25] which aims at recovering shape from a single image by describing the interference between light, geometry and reflectance of the object. It is impossible to unambiguously infer geometry, even if reflectance and lighting is known. One way to disambiguate shape-from-shading is to use multiple images instead of a single observation, the clue is to capture each image under different illumination. This approach is called photometric stereo [56] and tries to estimate the objects geometry and possibly its reflectance. One differs between calibrated PS and uncalibrated PS, where former assumes proper knowledge of the lighting conditions in the scene [45], and the latter, more complicated and ill-posed

case, tries to solve PS without lighting information [6, 8]. Besides the complexity of uncalibrated PS in lighting, it furthermore tends to assume a static camera-to-object relation, i.e. illumination is the only quantity changing in the scene. This makes PS impractical for applications where an object is captured under camera motion. Multi-view photometric stereo under unconstrained camera motion and illumination was presented in [5, 52], yet one needs example objects in order to recover reflectance properties. Other approaches need additional lighting calibration using a specular sphere during the capturing process and additionally capture multiple images from the same viewpoint under different illumination before moving the camera and repeating this procedure [61]. One can circumvent this by attaching a light source to the camera which then needs to be calibrated beforehand along with the extrinsic parameters of the sensor [24, 50]. Placing the object on a turntable and capture it from different viewpoints by rotating it implies different illumination on the object as well, even if lighting is static with respect to the camera [23, 62]. Although this allows for a non-static camera-to-object setup, the motion of the object relative to the camera remains limited due to the turntable setup. Structure-from-motion (SfM) [54] can be used to estimate odometry and an initial geometry [60], yet SfM relies on the photo consistency assumption which states that the same 3D point in the scene results in the same intensity across images taken from different viewpoints. When applying PS this assumption is heavily violated, which makes it difficult to apply a SfM approach out-of-the-box. A similar approach was considered in [36]. This is realized using LEDs soldered to a circuit board which is attached to a camera. The camera captures multiple images under different illumination, before repeating this process from different viewpoints. As the relative position of the LEDs to the camera is fixed and calibrated offline this approach performs calibrated PS. The advantage of having a sensor based depth estimate from an RGB-D camera seems to resolve the ambiguity arising in the uncalibrated PS case [44], which gives rise to use a low-resolution depth map. Only two approaches use RGB-D cameras in combination with multi-view PS [10, 62]. The capturing process in [10] is closely related to [36, 61], and [62] make use of a turntable. Thus, [10, 62] remain limited to their specific setup.

To the best of our knowledge, we are the first ones who not only combine both, depth super-resolution and multi-view photometric stereo, but also aim at estimating camera poses, illumination, reflectance and high resolution geometry jointly, thus circumventing the need of tedious offline calibration of a freely moving camera or lighting. As we will see later, our setup can be used out-of-the-box, i.e. we attach an LED light source to an RGB-D camera to enforce different illumination across different views while moving the camera.

3. Variational Depth Super-Resolution and Multi-View Photometric Stereo

Our proposed model consists of three major parts and we will discuss each in more mathematical detail before putting all together to an end-to-end variational formulation. We start with some theoretical foundation on depth super-resolution, before formalizing an image formation model for photometric stereo. An extension to multi-view PS is introduced using rigid body motions and how to naturally incorporate them to the classical PS model.

3.1. Depth Super-Resolution

Depth super-resolution tries to infer a high resolution (HR) depth map $z : \Omega_{\text{HR}} \subset \mathbb{R}^2 \rightarrow \mathbb{R}$, from $n \geq 1$ low-resolution (LR) depth measurements $z_i : \Omega_{\text{LR}} \subset \mathbb{R}^2 \rightarrow \mathbb{R}$, $i \in \mathcal{I} := \{0, \dots, n-1\}$. Formally, the corresponding forward process can be written as

$$z_i = D_i z + \varepsilon_{z_i}, \quad \forall i \in \mathcal{I}, \quad (1)$$

where $D_i : \Omega_{\text{HR}} \rightarrow \Omega_{\text{LR}}$ is some linear downsampling operator [55] and $\varepsilon_{z_i}(\mathbf{p}) \sim \mathcal{N}(0, \sigma_z^2)$, $\mathbf{p} \in \Omega_{\text{LR}}$, the realization of a stochastic process following a Gaussian distribution with mean 0 and variance σ_z^2 , see [30] for more details. If the cameras motion is not static, then these n LR depth measurements are captured from different viewpoints and usually one needs to define a so called reference frame. Furthermore, notion of the camera movement needs to be taken into account in the process described in (1). This can be realized in the linear operator D_i , such that not solely downsampling, but also image warping is considered [55]. Although we aim to tackle the problem depth super-resolution from multiple viewpoints, we will see later that our final variational model only depends on the LR reference depth, i.e. for a LR depth map we rely on a single observation (“ $n = 1$ ”) and denote it with z_0 (we also drop the dependency on i in D_i). The multi-view aspect in our model arises from the photometric constraint which we discuss in Section 3.3. To invert (1) we follow a bayesian rationale for our variational approach [41], i.e. (1) can be formulated as a minimization problem of the form

$$\min_{z: \Omega_{\text{HR}} \rightarrow \mathbb{R}} \tau \|Dz - z_0\|_{\ell_2(\Omega_{\text{LR}})}^2 + \mathcal{R}(z), \quad (2)$$

where $\|\cdot\|_{\ell_2(\Omega_{\text{LR}})}^2$ is the euclidean 2-norm over the image domain Ω_{LR} and τ is a trade-off parameter determining the impact of each of the two summands in (2). The so called regularization term $\mathcal{R}(z)$ is often based on prior knowledge. One choice is a total variation based regularization term incorporating sparse features of the intensity image [15]. Recently, a promising direction towards depth refinement is to use not only sparse, but dense information from the RGB image, i.e. use photometric cues to perform depth refinement [19, 42].

3.2. Photometric Stereo

Let us consider the following image formation model under Lambertian reflectance to establish a dense relation between shape and intensity:

$$I_c(\mathbf{p}) = \rho_c(\mathbf{p}) \int_{\mathbb{S}^2} \phi(\boldsymbol{\omega}) \max(0, \langle \boldsymbol{\omega}, \mathbf{n}(\mathbf{p}) \rangle) d\boldsymbol{\omega}, \quad (3)$$

where $I_c(\mathbf{p}), \rho_c(\mathbf{p}) \in \mathbb{R}$ is the image intensity, albedo in channel $c \in C = \{R, G, B\}$ at pixel $\mathbf{p} \in \Omega_{\text{HR}}$, respectively. The integral is referred to as shading and is carried out over the 3-dimensional unit sphere \mathbb{S}^2 , i.e. all possible lighting directions $\boldsymbol{\omega}$ with their corresponding intensity $\phi(\boldsymbol{\omega})$. $\mathbf{n}(\mathbf{p})$ is the surface normal represented at pixel $\mathbf{p} \in \Omega_{\text{HR}}$ and $\langle \cdot, \cdot \rangle : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}$ is the 3-dimensional dot product. Using the well-known first-order spherical harmonic approximation which accounts for 87% of real-world lighting [7], we can write (3) as

$$I_c(\mathbf{p}) = \rho_c(\mathbf{p}) \langle \mathbf{l}, \mathbf{m}(\mathbf{n}(\mathbf{p})) \rangle, \quad (4)$$

where the shading term is simplified to a 4-dimensional dot product with $\mathbf{l} \in \mathbb{R}^4$ representing the global lighting of the scene and $\mathbf{m}(\mathbf{n}(\mathbf{p})) \in \mathbb{R}^4$ are the spherical harmonic basis functions at pixel $\mathbf{p} \in \Omega_{\text{HR}}$,

$$\mathbf{m}(\mathbf{n}(\mathbf{p})) := [1, \mathbf{n}(\mathbf{p})^\top]^\top. \quad (5)$$

As our approach is running on RGB images, we stack the values for each channel to have

$$\mathbf{I}(\mathbf{p}) = \boldsymbol{\rho}(\mathbf{p}) \langle \mathbf{l}, \mathbf{m}(\mathbf{n}(\mathbf{p})) \rangle, \quad (6)$$

with $\mathbf{I}(\mathbf{p}) = [I_R(\mathbf{p}), I_G(\mathbf{p}), I_B(\mathbf{p})]^\top$ being the RGB intensity values and $\boldsymbol{\rho}(\mathbf{p}) = [\rho_R(\mathbf{p}), \rho_G(\mathbf{p}), \rho_B(\mathbf{p})]^\top$ being the RGB albedo values. We have now established a mathematical relationship between the color image \mathbf{I} and the geometry of the scene described with the surface normals \mathbf{n} . As we use (perspective) depth maps from an RGB-D sensor and we need to link absolute depth (, i.e. depth maps) with relative geometric information (, i.e. surface normals), we parametrize the normal-field \mathbf{n} in terms of depth z under perspective projection and assuming differentiability [46],

$$\mathbf{n}(z(\mathbf{p})) = \frac{1}{d\mathcal{A}(z(\mathbf{p}))} \left[-1 - \langle \mathbf{p} - \mathbf{c}, \nabla z(\mathbf{p}) \rangle \right] \in \mathbb{S}^2, \quad (7)$$

where $d\mathcal{A}(z(\mathbf{p})) = \sqrt{|f\nabla z(\mathbf{p})|^2 + (1 + \langle \mathbf{p} - \mathbf{c}, \nabla z(\mathbf{p}) \rangle)^2}$ is the surface area element under perspective projection [17], ∇ is the gradient operator and $|\cdot|$ is the 2-dimensional euclidean norm, f and $\mathbf{c} = [c_x, c_y]^\top$ are camera parameters, i.e. the focal length and principal point, respectively. To avoid clutter in notation we drop the argument dependency on the pixel \mathbf{p} and rather consider each quantity on the whole image domain Ω_{HR} ,

$$\mathbf{I} = \boldsymbol{\rho} \langle \mathbf{l}, \mathbf{m}[\mathbf{n}[z]] \rangle, \quad (8)$$

with the RGB intensity image $I : \Omega_{\text{HR}} \rightarrow \mathbb{R}^3$, the albedo $\rho : \Omega_{\text{HR}} \rightarrow \mathbb{R}^3$, the global lighting vector $l \in \mathbb{R}^4$, and the stacked spherical harmonic basis functions m depending on the surface normals n parametrized with z , i.e. we actually consider a vector-field $m[n[z]] : \Omega_{\text{HR}} \rightarrow \mathbb{R}^4$. In order to perform photometric stereo, $n \geq 3$ images under different illumination need to be acquired [56], and since geometry and reflectance are properties of the scene and stay fixed, these changes become only apparent in lighting and intensity, this can be mathematically realized with

$$I_i = \rho \langle l_i, m[n[z]] \rangle, \quad i \in \mathcal{I}. \quad (9)$$

Practically, changes in illumination can be achieved by using a light stage, illuminating the scene with a light-bulb or LED light source, or capturing pictures during different times of the day or even during different seasons of the year [3, 4, 28].

Based on (2) and (9), it was shown in [19, 44] that under the assumption of $\{I_i\}_{i \in \mathcal{I}}$ being corrupted with homoskedastically Gaussian-distributed noise a purely data-driven, photometric regularization term can be established, $\mathcal{R}_{[44]}(z) = \sum_{i \in \mathcal{I}} \|I_i - \rho \langle l_i, m[n[z]] \rangle\|_{\ell_2(\Omega_{\text{HR}})}^2$. Apart from avoiding a man-made regularization term, both, the depth prior term (left summand in (2)) and photometric term $\mathcal{R}_{[44]}(z)$ can be each others helping hand. Treated separately, depth super-resolution and uncalibrated photometric stereo are ill-posed problems by itself [6, 8, 14], but the combination of both can help disambiguating the other: The depth prior term resulting from an RGB-D sensors initial depth uniquely defines the Lorentz transformation in order to disambiguate uncalibrated photometric stereo [9]. Similarly, using the information comprised in the intensity image and thus in the photometric term can help to result in fine detailed depth estimates and to disambiguate depth super-resolution. To this end it seems intuitive to combine both approaches to get the best out of both worlds. Yet, [44] assumes a static camera-to-object setup during the capturing process which makes it impractical to use in a scenario where camera motion is present. A natural way to capture data would be to attach the light source to the RGB-D sensor and freely move the camera around. This results in an easy way to capture data from different viewpoints while simultaneously changing illumination. Furthermore, resorting to RGB-D cameras has the advantage of robustly capturing depth data, independently of illumination changes in the visible spectrum, which gives coarse, but robust knowledge of the geometry of the scene which can be used during inference. The next section is concerned with how to incorporate the practical setup of attaching an LED light source to an RGB-D camera in the variational formulation. To this end, rigid body motions and how we will use them to model multi-view photometric stereo are discussed.

3.3. Using Rigid Body Motion for Multi-View Photometric Stereo

Rigid body motions can be used to describe how a camera moves in space, and additionally gives a relation between pixel positions of different images captured from different viewpoints. Under perspective projection a pixel $p \in \Omega_{\text{HR}} \subset \mathbb{R}^2$ and its depth value $z(p)$ is used to express its conjugate 3D point $P \in \mathbb{R}^3$ in the scene,

$$\pi^{-1}(p, z(p)) = z(p) \begin{bmatrix} f^{-1}(p - c) \\ 1 \end{bmatrix} = P, \quad (10)$$

where $\pi^{-1} : \Omega_{\text{HR}} \times \mathbb{R} \rightarrow \mathbb{R}^3$ can be referred to as the reprojection function. The relation between P seen from one camera and the same corresponding point \tilde{P} seen from a second camera can be expressed as a rigid body motion. This comprises a rotation $\mathbf{R} \in SO(3)$, where $SO(3)$ is the space of 3×3 dimensional rotation matrices, and a translation $\mathbf{t} \in \mathbb{R}^3$,

$$\mathbf{T}[\mathbf{R}, \mathbf{t}](P) = \mathbf{R}P + \mathbf{t} = \tilde{P}, \quad (11)$$

with $\mathbf{T}[\mathbf{R}, \mathbf{t}] : \mathbb{R}^3 \rightarrow \mathbb{R}^3$. After applying the coordinate transformation with $\mathbf{T}[\mathbf{R}, \mathbf{t}]$, the resulting 3D point $\tilde{P} = (\tilde{X}, \tilde{Y}, \tilde{Z})^\top$ has to be projected on the image plane again,

$$\pi(\tilde{P}) = \frac{f}{\tilde{Z}} \begin{bmatrix} \tilde{X} \\ \tilde{Y} \end{bmatrix} - c = \tilde{p}, \quad (12)$$

where $\pi : \mathbb{R}^3 \rightarrow \Omega_{\text{HR}}$. We can describe the process of re-projection (10), camera coordinate transformation (11) and projection (12) using a single warping function $w[\mathbf{R}, \mathbf{t}, z] : \Omega_{\text{HR}} \rightarrow \Omega_{\text{HR}}$,

$$w[\mathbf{R}, \mathbf{t}, z](p) = \pi(\mathbf{T}[\mathbf{R}, \mathbf{t}](\pi^{-1}(p, z(p)))) \quad (13)$$

This warping function describes how a pixel from one image frame is transformed to locate the same 3D point in the second image. As we have multiple (namely $n \geq 3$) observations due to the multi-view setup, we have to keep track of each rotation and translation. To this end, we transform the camera coordinate system from the reference frame (0-th frame, cf. Section 3.1) to the i -th frame using the rotation \mathbf{R}_i and translation \mathbf{t}_i . Note that this transformation needs only the depth map from the reference frame, as the re-projection is applied to the reference frame, thus only the rigid body motion is depending on the camera pose of the i -th frame. Photometrically, under the Lambertian assumption the intensity in the i -th frame at pixel \tilde{p}_i can be described as illuminating the scene from the i -th frame, but seen from the reference frame at pixel p , mathematically this can be written by combining (9) and (13) as

$$I_i \circ w[\mathbf{R}_i, \mathbf{t}_i, z] = \rho \langle l_i, m[n[z]] \rangle, \quad i \in \mathcal{I}. \quad (14)$$

For simplicity, we define $\tilde{I}_i[\mathbf{R}_i, \mathbf{t}_i, z] := I_i \circ w[\mathbf{R}_i, \mathbf{t}_i, z]$ and use this representation throughout the rest of this work, before discussing our proposed variational setup.

3.4. Proposed Variational Formulation

Putting together the forward model on depth super-resolution (1) and multi-view uncalibrated photometric stereo (14), we formulate the joint recovery of camera motion, geometry, reflectance and lighting as the following variational problem

$$\min_{\substack{\{\mathbf{R}_i\} \in SO(3) \\ \{\mathbf{t}_i\} \in \mathbb{R}^3 \\ z: \Omega_{\text{HR}} \rightarrow \mathbb{R} \\ \rho: \Omega_{\text{HR}} \rightarrow \mathbb{R}^3 \\ \{\mathbf{l}_i\} \in \mathbb{R}^4}} \sum_{i \in \mathcal{I}} \Phi_{\mathcal{C}} \left(\tilde{\mathbf{I}}_i [\mathbf{R}_i, \mathbf{t}_i, z] - \rho \langle \mathbf{m} [\mathbf{n} [z]], \mathbf{l}_i \rangle \right) + \tau \|Dz - z_0\|_{\ell_2(\Omega_{\text{LR}})}^2. \quad (15)$$

We denote with $\Phi_{\mathcal{C}}(\cdot)$ the sum over all pixels and all channels over data-fitting discrepancy,

$$\Phi_{\mathcal{C}} \left(\tilde{\mathbf{I}}_i [\mathbf{R}_i, \mathbf{t}_i, z] - \rho \langle \mathbf{m} [\mathbf{n} [z]], \mathbf{l}_i \rangle \right) = \sum_{\mathbf{p} \in \Omega_{\text{HR}} \subset \mathcal{C}} \sum_{c \in \mathcal{C}} \phi_{\mathcal{C}} \left(\tilde{\mathbf{I}}_{i,c} [\mathbf{R}_i, \mathbf{t}_i, z](\mathbf{p}) - \rho_c(\mathbf{p}) \langle \mathbf{m} [\mathbf{n} [z]](\mathbf{p}), \mathbf{l}_i \rangle \right), \quad (16)$$

where $\phi_{\mathcal{C}}(r) = \frac{\lambda^2}{2} \log \left(1 + \frac{r^2}{\lambda^2} \right)$ is Cauchy's M-estimator (evaluated at residual r) to make the approach robust against outliers. In the context of photometric stereo outliers can be self- or cast-shadows, specularities, and image noise or quantization, see [47] for more detail. The parameter λ is a tunable hyper-parameter and we heuristically found $\lambda = 0.04$ to be a good choice. We will now discuss our numerical approach to solve the variational problem on depth super-resolution using multi-view uncalibrated PS.

4. Numerical Resolution

The problem depicted in (15) is nonconvex wrt. z (due to the normalization factor, cf. (7)), \mathbf{R}_i , as well as Cauchy's M-estimator. This makes our variational problem difficult to solve and we will now discuss a strategy to make optimization tractable. To this end, we follow [22, 44] and rewrite the normals in terms of its *unnormalized normal* $\tilde{\mathbf{n}} [z]$ and its norm, $\mathbf{n} [z] = \frac{\tilde{\mathbf{n}} [z]}{|\tilde{\mathbf{n}} [z]|} = \frac{\tilde{\mathbf{n}} [z]}{d\mathcal{A} [z]}$, this allows us to split the nonconvexity induced by the normal into its convex (linear) and its nonconvex part. Camera rotation and translation represent elements of the Lie group $SE(3)$, which for the numerical solution we parameterize by the corresponding Lie algebra element, i.e. the 6-dimensional twist coordinates $\xi_i \in \mathbb{R}^6$, $\mathbf{R}_i = \mathbf{R}_i(\xi_i)$ and $\mathbf{t}_i = \mathbf{t}_i(\xi_i)$. This allows to optimize over the Euclidean parameters ξ_i , cf. [29, 39]. We can now rewrite (15) as

$$\min_{\substack{\{\xi_i\} \in \mathbb{R}^6 \\ z: \Omega_{\text{HR}} \rightarrow \mathbb{R} \\ \rho: \Omega_{\text{HR}} \rightarrow \mathbb{R}^3 \\ \{\mathbf{l}_i\} \in \mathbb{R}^4}} \sum_{i \in \mathcal{I}} \Phi_{\mathcal{C}} \left(\tilde{\mathbf{I}}_i [\mathbf{R}_i(\xi_i), \mathbf{t}_i(\xi_i), z] - \rho \left\langle \mathbf{m} \left[\frac{\tilde{\mathbf{n}} [z]}{d\mathcal{A} [z]} \right], \mathbf{l}_i \right\rangle \right) + \tau \|Dz - z_0\|_{\ell_2(\Omega_{\text{LR}})}^2. \quad (17)$$

The nonconvexity induced by $\phi_{\mathcal{C}}$ in (16) gives rise to perform a reweighted least squares optimization scheme, cf. [47]. This boils down to approximating $\phi_{\mathcal{C}}$ at the k -th iteration and residual $r^{(k)}$ with its weighted least squares estimator, $\phi_{\mathcal{C}}(r^{(k)}) \approx \phi_{\mathcal{C}}^{(k)}(r^{(k)}) := w^{(k)}(r^{(k)}) \phi_2(r^{(k)})$, where $w^{(k)}(r^{(k)}) = \frac{\phi_{\mathcal{C}}'(r^{(k)})}{r^{(k)}}$ and $\phi_2(r^{(k)}) = |r^{(k)}|^2$. Analogously, we denote this approximation at the k -th iteration with $\Phi_{\mathcal{C}}^{(k)}$. That is, at the (k) -th iteration with the given estimates ($\{\xi_i^{(k)}\}, z^{(k)}, \rho^{(k)}, \{\mathbf{l}_i^{(k)}\}$) we update each quantity according to the following sweep:

$$\xi_i^{(k+1)} = \operatorname{argmin}_{\xi_i \in \mathbb{R}^6} \Phi_{\mathcal{C}}^{(k)} \left(\tilde{\mathbf{I}}_i [\mathbf{R}_i(\xi_i), \mathbf{t}_i(\xi_i), z^{(k)}] \right) \quad (18)$$

$$- \rho^{(k)} \left\langle \mathbf{m} \left[\frac{\tilde{\mathbf{n}} [z^{(k)}]}{d\mathcal{A} [z^{(k)}]} \right], \mathbf{l}_i^{(k)} \right\rangle, \quad \forall i \in \mathcal{I},$$

$$z^{(k+1)} = \operatorname{argmin}_{z: \Omega_{\text{HR}} \rightarrow \mathbb{R}} \sum_{i \in \mathcal{I}} \Phi_{\mathcal{C}}^{(k+\frac{1}{4})} \left(\tilde{\mathbf{I}}_i [\mathbf{R}_i(\xi_i^{(k+1)}), \mathbf{t}_i(\xi_i^{(k+1)}), z^{(k)}] \right) - \rho^{(k)} \left\langle \mathbf{m} \left[\frac{\tilde{\mathbf{n}} [z]}{d\mathcal{A} [z^{(k)}]} \right], \mathbf{l}_i^{(k)} \right\rangle + \tau \|Dz - z_0\|_{\ell_2(\Omega_{\text{LR}})}^2, \quad (19)$$

$$\rho^{(k+1)} = \operatorname{argmin}_{\rho: \Omega_{\text{HR}} \rightarrow \mathbb{R}^3} \sum_{i \in \mathcal{I}} \Phi_{\mathcal{C}}^{(k+\frac{2}{4})} \left(\tilde{\mathbf{I}}_i [\mathbf{R}_i(\xi_i^{(k+1)}), \mathbf{t}_i(\xi_i^{(k+1)}), z^{(k+1)}] \right) - \rho \left\langle \mathbf{m} \left[\frac{\tilde{\mathbf{n}} [z^{(k+1)}]}{d\mathcal{A} [z^{(k+1)}]} \right], \mathbf{l}_i^{(k)} \right\rangle, \quad (20)$$

$$\mathbf{l}_i^{(k+1)} = \operatorname{argmin}_{\mathbf{l}_i \in \mathbb{R}^4} \Phi_{\mathcal{C}}^{(k+\frac{3}{4})} \left(\tilde{\mathbf{I}}_i [\mathbf{R}_i(\xi_i^{(k+1)}), \mathbf{t}_i(\xi_i^{(k+1)}), z^{(k+1)}] \right) - \rho^{(k+1)} \left\langle \mathbf{m} \left[\frac{\tilde{\mathbf{n}} [z^{(k+1)}]}{d\mathcal{A} [z^{(k+1)}]} \right], \mathbf{l}_i \right\rangle, \quad \forall i \in \mathcal{I}. \quad (21)$$

We follow a coarse-to-fine strategy with five pyramid levels, i.e. on the coarsest level the resolution is $\frac{1}{16}$ of the original size. For each pyramid level we solve (17) using (18) – (21). The estimated quantities are then (bilinearly up-sampled for z and ρ) used as input for the next pyramid level. We initialize $\xi_0^{(0)} \equiv \mathbf{0}$ (identity), and $\xi_{i+1}^{(0)} = \xi_i^{(1)}$, $i \in \mathcal{I}$ (small baseline assumption). $z^{(0)}$ and $\rho^{(0)}$ are bilinearly downsampled versions of z_0 and \mathbf{I}_0 , respectively and $\mathbf{l}_i^{(0)} = (0.2, 0, 0, -1)$ (little ambient amount and frontal lighting), $\forall i \in \mathcal{I}$. Our algorithm has converged if the relative energy falls below a value of 10^{-5} , which on average took 20 iterations for each pyramid level. Note that (19) – (21) are linear least square problems, while we resort to Gauss-Newton iterations for (18), cf. [29]. Next, we quantitatively and qualitatively evaluate our approach on synthetic and real-world data.

5. Empirical Validation

We validate our method on synthetic and real-world data. We draft insightful experiments of the involved parameters and compare against other state-of-the-art depth SR approaches.

5.1. Synthetic Data

We use two publicly available datasets for synthetic evaluation, ‘‘Joyful Yell’’ [1] and ‘‘Stanford Bunny’’ [53]. Both objects are used and colored within the rendering software Blender. A virtual camera with collocated light source is used to render images from different viewpoints. The resulting data, with its corresponding ground truth is then used for evaluation. Low-resolution depth maps ($\frac{1}{2}, \frac{1}{4}, \frac{1}{8}$ of original size) are generated by adding non-uniform zero-mean Gaussian noise with standard deviation 10^{-5} times the squared original depth value [30], cf. Figure 2 for an illustration.

While having ground-truth at hand using synthetic data, we evaluate how n and τ impacts the mean angular error (MAE) and the root mean squared error (RMSE).

Parameter Evaluation. To evaluate for n , we use the first $n \in \{5, 10, 15, 20, 25, 30\}$ frames of the image sequence. The results obtained with our proposed approach are reported in Figure 3. With increasing number of images the error in both error metrics is decreasing. Yet, for more than 15–20 images the error increases again. This is most likely due to a larger baseline between the images and the 0-th frame, error is accumulated in camera pose estimation and thus depth inference suffers after warping. Given this insight, we choose $n = 20$ for all our experiments. Next, we evaluate against τ . Due to the pyramid scheme and to be more robust across datasets (larger/smaller objects, different depth magnitude, etc.) we perform parameter normalization, $\tau = \frac{n \text{mean}(\{I_i\})^2 |\Omega_{HR}| |C|}{\text{mean}(z_0)^2 |\Omega_{LR}|} \tilde{\tau}$, where $\text{mean}(\cdot)$ is the mean operator, across all input RGB images or the low-resolution depth, $|\cdot|$ here refers to the cardinality of Ω_{HR} , C and Ω_{LR} and $\tilde{\tau}$ is the parameter we evaluate against. Results are shown in Figure 4. For low values of $\tilde{\tau}$ the depth prior term can be considered to have little to no influence and the approach can be compared to an uncalibrated PS one. Thus, the overall shape (RMSE is large) suffers from the ill-posedness of uncalibrated PS [6]. On the other hand, large values of $\tilde{\tau}$ result in a pure depth super-resolution approach, and the result can only get as good as the input depth, i.e. missing fine geometric information. This is resembled in the MAE value, which heavily increases for too large values of $\tilde{\tau}$. Given this insight, $\tilde{\tau} \in [1, 100]$ seems a good trade-off between fine-scale details (low MAE) and global geometry fit (low RMSE), and we set $\tilde{\tau} = 10$ in all further experiments. Note that even for scaling-factor (SF) 8 the difference in MAE and RMSE is small compared SF 4 or 2, showing that our approach is robust across different SF.

Dataset	Metric	[35]	[44]	[33]	[21]	Ours
Joyful Yell	MAE	41.43	12.73	26.32	27.42	12.54
	RMSE	0.062	0.062	0.450	0.099	0.039
Stanford Bunny	MAE	32.33	21.42	20.89	25.86	16.06
	RMSE	0.028	0.028	0.114	0.044	0.022

Table 1: Quantitative results on the Joyful Yell and Stanford Bunny of our proposed method against other state-of-the-art depth super-resolution approaches.

Comparison with other Methods. Having the parameters fixed, we quantitatively and qualitatively evaluate our approach against other state-of-the-art depth super-resolution approaches. We compare to a mathematically principled robust optimization framework for color guided depth map restoration [35]. To show the impact of jointly solving for camera pose and shape, reflectance and lighting, we compare against [44], for which we first solve for camera poses (basically solving only (18) with our framework by setting $z = z_0$ and $\rho \langle m[n[z]], l_i \rangle = I_0$) followed by applying the model of [44] to the warped images. In [33] a unified approach to a hierarchical optimization framework and fast global smoothing performing cascaded interpolation with alternating RGB guidance is used. To address the impact of the multi-view setup, we also compare to a single-shot depth super-resolution scheme incorporating photometric cues using the shape-from-shading principle [21]. Table 1 shows quantitative (on Joyful Yell and Stanford Bunny) and Figure 6 shows qualitative (on Joyful Yell, Stanford Bunny and real-world data we captured ourselves) results. In numbers, our approach performs best in terms of both error metrics, showing the superiority compared to other state-of-the-art methods. This can also be verified visually, while the results of [33] are too smooth, the ones of [35] appear too noisy. Better shape estimates can be achieved with [21, 44], showing the positive impact on photometric approaches for depth super-resolution. Next, we discuss our hardware setup and conclude the experimental section with a comparison on real-world results.

5.2. Setup and Real-World Data

Real-world data is captured using an Intel RealSense D415 RGB-D camera with attached LED light source, cf. Figure 5. The resolution of RGB and depth is 1280×720 and 640×360 , respectively. We captured $n = 20$ images of each object shown in Figure 6 by moving the camera at a distance of 0.5–2m under low ambient light such that the LED illumination dominates. Although automatic approaches exist [20], we manually segment each image, and the first frame defines the reference image I_0 and depth z_0 . No further calibration is needed while f and c can be accessed using Intels RealSense camera software [2].

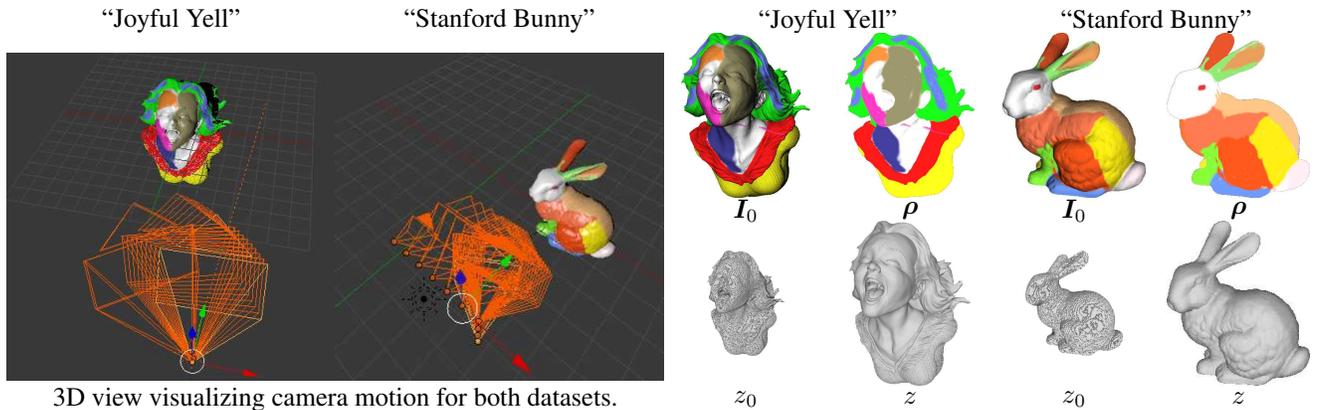


Figure 2: (Left) Visualization of “Joyful Yell” and “Stanford Bunny” in space with corresponding camera motion used to render both image sequences. The motion for the “Joyful Yell” dataset comprises rotation only, the “Stanford Bunny” got captured under rotation and translation. The experiments will show that our approach handles both camera motions with similar ease and precision (quite in contrast to passive sensing approaches where depth cannot be recovered under pure rotation). (Right) The rendered reference image, ground-truth albedo, low-resolution reference depth and ground-truth depth.

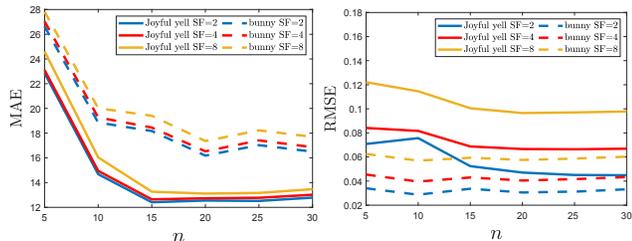


Figure 3: Impact of the number of input images n on the accuracy for three scaling factors (SF). Left is the Mean Angular Error (MAE, in degrees) on the normals. Right is the Root Mean Square Error (RMSE, in arbitrary units) on depth. 15 – 20 images are enough to obtain accurate results. More than 20 images seem to increase the error, due to a larger baseline to the reference frame.

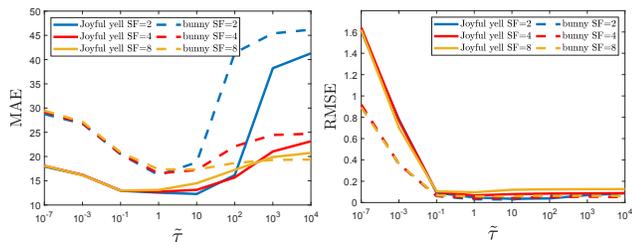


Figure 4: Impact of the input parameter $\tilde{\tau}$ on the accuracy for three scaling factors (SF) under two different datasets. Left is the Mean Angular Error (MAE, in degrees) on the normals. Right is the Root Mean Square Error (RMSE, in arbitrary units) on depth. The interval $\tau \in [1, 100]$ leads globally accurate results (low RMSE) with fine geometric details (low MAE).



Figure 5: Our setup we used to capture real-world data. An LED light source is attached to an Intel RealSense D415 RGB-D camera. We then move the camera to capture the objects of interest, see Figure 6. Different illumination is induced by camera motion and the attached LED.

The qualitative results can be seen in Figure 6. While too smooth depth estimates are recovered for [33], better results under the appearance of some remaining noise (Joyful Yell and Hat) are achieved by [35]. The importance of joint estimation of camera motion and geometry, albedo and lighting becomes apparent when comparing our results to the ones of [44]. Although, the results of [44] are satisfactory on the Jacket and Hat dataset, the approach is not able to recover the fine details of the stitching on the Basecap. These fine details are recovered by [21], but artifacts are visible on Jacket and Pillow. Our approach gives satisfactory results across all datasets, showing that multi-view uncalibrated photometric stereo is indeed helpful to recover high resolution fine detailed depth.

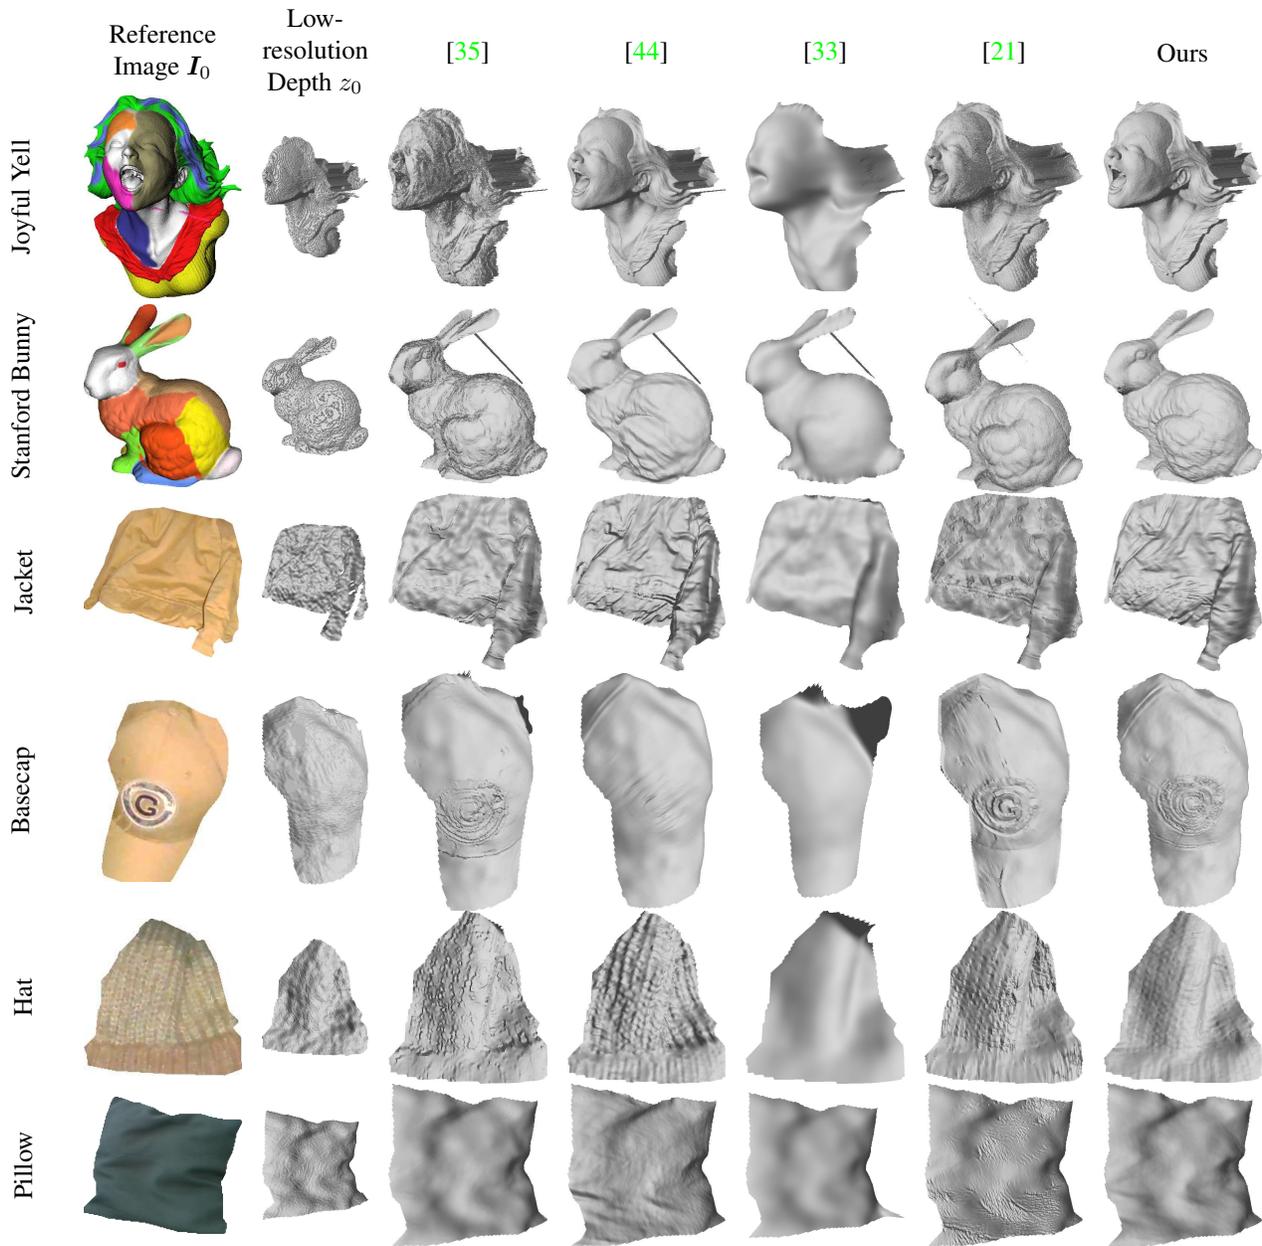


Figure 6: Results of state-of-the-art methods and our approach on challenging synthetic (Joyful Yell and Stanford Bunny) and real-world (rest) datasets. Our approach is the only one giving satisfactory results across all datasets.

6. Conclusion

We presented a novel approach to depth super-resolution using multi-view uncalibrated photometric stereo using a simple setup by attaching an LED light source to an RGB-D camera. No calibration on lighting and camera motion is necessary, due to our end-to-end variational approach which jointly optimizes over camera motion, geometry, reflectance and lighting. We demonstrated on multiple challenging synthetic and real-world datasets the superiority of our method

compared to other state-of-the-art approaches.

In the future, aiming at modelling the LED as a point light source or using all available depth images from the RGB-D sensor seem to be promising directions in order to further improve geometry estimation.

7. Acknowledgment

This work was supported by the ERC Consolidator Grant *3D Reloaded*.

References

- [1] The Joyful Yell, 2015. <https://www.thingiverse.com/thing:897412>. 6
- [2] Intel RealSense SDK 2.0, 2019. <https://github.com/IntelRealSense/librealsense>. 6
- [3] A. Abrams, C. Hawley, and R. Pless. Heliometric stereo: Shape from sun position. In *European Conference on Computer Vision*, pages 357–370, 2012. 4
- [4] J. Ackermann, F. Langguth, S. Fuhrmann, and M. Goesele. Photometric stereo for outdoor webcams. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 262–269, 2012. 4
- [5] J. Ackermann, F. Langguth, S. Fuhrmann, A. Kuijper, and M. Goesele. Multi-view photometric stereo by example. In *2014 International Conference on 3D Vision*, pages 259–266, 2014. 2
- [6] R. Basri, D. Jacobs, and I. Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of Computer Vision*, 72(3):239–257, 2007. 2, 4, 6
- [7] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(2):218–233, 2003. 3
- [8] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille. The bas-relief ambiguity. *International Journal of Computer Vision*, 35(1):33–44, 1999. 2, 4
- [9] M. Brahimi, Y. Quéau, B. Haefner, and D. Cremers. On well-posedness of uncalibrated photometric stereo under general lighting. *arXiv preprint arXiv:1911.07268*, 2019. 4
- [10] E. Bylow, R. Maier, F. Kahl, and C. Olsson. Combining depth fusion and photometric stereo for fine-detailed 3d models. In *Scandinavian Conference on Image Analysis*, pages 261–274, 2019. 2
- [11] M. Cristani, D. S. Cheng, V. Murino, and D. Pannullo. Distilling information with super-resolution for video surveillance. In *Proceedings of the ACM 2Nd International Workshop on Video Surveillance & Sensor Networks*, pages 2–11, 2004. 2
- [12] C. Dong, C. C. Loy, K. He, and X. Tang. Learning a deep convolutional network for image super-resolution. In *European conference on computer vision*, pages 184–199, 2014. 2
- [13] C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 38(2):295–307, 2015. 2
- [14] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar. Advances and challenges in super-resolution. *International Journal of Imaging Systems and Technology*, 14(2):47–57, 2004. 4
- [15] D. Ferstl, C. Reinbacher, R. Ranftl, M. Rütther, and H. Bischof. Image guided depth upsampling using anisotropic total generalized variation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 993–1000, 2013. 2, 3
- [16] B. Goldluecke, M. Aubry, K. Kolev, and D. Cremers. A super-resolution framework for high-accuracy multiview reconstruction. *International Journal of Computer Vision*, 106:172–191, 2014. 2
- [17] G. Graber, J. Balzer, S. Soatto, and T. Pock. Efficient minimal-surface regularization of perspective depth maps in variational stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 511–520, 2015. 3
- [18] H. Greenspan. Super-resolution in medical imaging. *The Computer Journal*, 52(1):43–63, 2008. 2
- [19] B. Haefner, S. Peng, A. Verma, Y. Quéau, and D. Cremers. Photometric depth super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019. 2, 3, 4
- [20] B. Haefner, Y. Quéau, and D. Cremers. Photometric segmentation: Simultaneous photometric stereo and masking. In *International Conference on 3D Vision*, 2019. 6
- [21] B. Haefner, Y. Quéau, T. Möllenhoff, and D. Cremers. Fight ill-posedness with ill-posedness: Single-shot variational depth super-resolution from shading. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2, 6, 7, 8
- [22] B. Haefner, Z. Ye, M. Gao, T. Wu, Y. Quéau, and D. Cremers. Variational uncalibrated photometric stereo under general lighting. In *International Conference on Computer Vision*, 2019. 5
- [23] C. Hernandez, G. Vogiatzis, and R. Cipolla. Multiview photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3):548–554, 2008. 2
- [24] T. Higo, Y. Matsushita, N. Joshi, and K. Ikeuchi. A handheld photometric stereo camera for 3-d modeling. In *2009 IEEE International Conference on Computer Vision*, pages 1234–1241, 2009. 2
- [25] B. K. P. Horn and M. J. Brooks, editors. *Shape from Shading*. MIT Press, 1989. 2
- [26] B. Huhle, T. Schairer, P. Jenke, and W. Straßer. Fusion of range and color images for denoising and resolution enhancement with a non-local filter. *Computer Vision and Image Understanding*, 114(12):1336–1345, 2010. 2
- [27] T.-W. Hui, C. C. Loy, and X. Tang. Depth map super-resolution by deep multi-scale guidance. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *European Conference on Computer Vision*, pages 353–369. Springer International Publishing, 2016. 2
- [28] J. Jung, J.-Y. Lee, and I. S. Kweon. One-day outdoor photometric stereo using skylight estimation. *International Journal of Computer Vision*, pages 1–17, 2019. 4
- [29] C. Kerl, J. Sturm, and D. Cremers. Robust odometry estimation for rgb-d cameras. In *International Conference on Robotics and Automation*, 2013. 5
- [30] K. Khoshelham and S. O. Elberink. Accuracy and resolution of kinect depth data for indoor mapping applications. *Sensors*, 12(2):1437–1454, 2012. 3, 6
- [31] M. Kiechle, S. Hawe, and M. Kleinsteuber. A joint intensity and depth co-sparse analysis model for depth map super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1545–1552, 2013. 2
- [32] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al.

- Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4681–4690, 2017. 2
- [33] Y. Li, D. Min, M. N. Do, and J. Lu. Fast guided global interpolation for depth and motion. In *European Conference on Computer Vision*, 2016. 6, 7, 8
- [34] Y. Li, T. Xue, L. Sun, and J. Liu. Joint example-based depth map super-resolution. In *2012 IEEE International Conference on Multimedia and Expo*, pages 152–157, 2012. 2
- [35] W. Liu, X. Chen, J. Yang, and Q. Wu. Robust color guided depth map restoration. *IEEE Transactions on Image Processing*, 26(1):315–327, 2017. 6, 7, 8
- [36] F. Logothetis, R. Mecca, and R. Cipolla. A differential volumetric approach to multi-view photometric stereo. *arXiv preprint arXiv:1811.01984*, 2018. 2
- [37] J. Lu, D. Min, R. S. Pahwa, and M. N. Do. A revisit to mrf-based depth map super-resolution and enhancement. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 985–988, 2011. 2
- [38] Z. Lu, Y.-W. Tai, F. Deng, M. Ben-Ezra, and M. S. Brown. A 3d imaging framework based on high-resolution photometric-stereo and low-resolution depth. *International Journal of Computer Vision*, 102(1-3):18–32, 2013. 2
- [39] Y. Ma, S. Soatto, J. Kosecka, and S. S. Sastry. *An invitation to 3-d vision: from images to geometric models*, volume 26. Springer Science & Business Media, 2012. 5
- [40] R. Maier, J. Stückler, and D. Cremers. Super-resolution keyframe fusion for 3d modeling with high-quality textures. In *2015 International Conference on 3D Vision*, pages 536–544, 2015. 2
- [41] D. Mumford. Bayesian rationale for the variational formulation. In *Geometry-driven Diffusion in Computer Vision*, pages 135–146. Springer, 1994. 3
- [42] R. Or-El, G. Rosman, A. Wetzler, R. Kimmel, and A. M. Bruckstein. Rgb-d-fusion: Real-time high precision depth recovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5407–5416, 2015. 3
- [43] J. Park, H. Kim, Y.-W. Tai, M. S. Brown, and I. S. Kweon. High-quality depth map upsampling and completion for rgb-d cameras. *IEEE Transactions on Image Processing*, 23(12):5559–5572, 2014. 2
- [44] S. Peng, B. Haefner, Y. Quéau, and D. Cremers. Depth super-resolution meets uncalibrated photometric stereo. In *International Conference on Computer Vision Workshops*, 2017. 2, 4, 5, 6, 7, 8
- [45] Y. Quéau, B. Durix, T. Wu, D. Cremers, F. Lauze, and J.-D. Durou. Led-based photometric stereo: Modeling, calibration and numerical solution. *Journal of Mathematical Imaging and Vision*, 07 2017. 2
- [46] Y. Quéau, J.-D. Durou, and J.-F. Aujol. Normal integration: a survey. *Journal of Mathematical Imaging and Vision*, 60(4):576–593, 2018. 3
- [47] Y. Quéau, T. Wu, F. Lauze, J.-D. Durou, and D. Cremers. A non-convex variational approach to photometric stereo under inaccurate lighting. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 5
- [48] G. Riegler, D. Ferstl, M. Rüther, and H. Bischof. A deep primal-dual network for guided depth super-resolution. *arXiv preprint arXiv:1607.08569*, 2016. 2
- [49] G. Riegler, M. Rüther, and H. Bischof. Atgv-net: Accurate depth super-resolution. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *European Conference on Computer Vision*, pages 268–284, 2016. 2
- [50] D. Simakov, D. Frolova, and R. Basri. Dense shape reconstruction of a moving object under arbitrary, unknown lighting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1202–1209, 2003. 2
- [51] X. Song, Y. Dai, and X. Qin. Deep depth super-resolution: Learning depth super-resolution using deep convolutional neural network. In S.-H. Lai, V. Lepetit, K. Nishino, and Y. Sato, editors, *Asian Conference on Computer Vision*, pages 360–376, 2017. 2
- [52] A. Treuille, A. Hertzmann, and S. M. Seitz. Example-based stereo with general brdfs. In *European Conference on Computer Vision*, pages 457–469, 2004. 2
- [53] G. Turk and M. Levoy. Zipped polygon meshes from range images. In *Proceedings of the 21st annual Conference on Computer Graphics and Interactive Techniques*, pages 311–318, 1994. 6
- [54] S. Ullman. The interpretation of structure from motion. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 203(1153):405–426, 1979. 2
- [55] M. Unger, T. Pock, M. Werlberger, and H. Bischof. A convex approach for variational super-resolution. In *Proceedings of the 32nd DAGM Conference on Pattern Recognition*, pages 313–322, 2010. 3
- [56] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139 – 144, 1980. 2, 4
- [57] J. Yang, J. Wright, T. Huang, and Y. Ma. Image super-resolution as sparse representation of raw image patches. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. 2
- [58] J. Yang, J. Wright, T. S. Huang, and Y. Ma. Image super-resolution via sparse representation. *IEEE Transactions on Image Processing*, 19(11):2861–2873, 2010. 2
- [59] Q. Yang, R. Yang, J. Davis, and D. Nistér. Spatial-depth super resolution for range images. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007. 2
- [60] L. Zhang et al. Shape and motion under varying illumination: Unifying structure from motion, photometric stereo, and multiview stereo. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 618–625, 2003. 2
- [61] Z. Zhou, Z. Wu, and P. Tan. Multi-view photometric stereo with spatially varying isotropic materials. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1482–1489, 2013. 2
- [62] X. Zuo, S. Wang, J. Zheng, and R. Yang. Detailed surface geometry and albedo recovery from rgb-d video under natural illumination. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3133–3142, 2017. 2