

# Tackling Shapes and BRDFs Head-on

Stamatios Georgoulis, Marc Proesmans and Luc Van Gool  
ESAT-PSI/VISICS, KU Leuven

stam.georgoulis, marc.proesmans, luc.vangool @esat.kuleuven.be

## Abstract

*In this work, we investigate the use of simple flash-based photography to capture an object’s 3D shape and reflectance characteristics at the same time. The presented method is based on the principles of Structure from Motion (SfM) and Photometric Stereo (PS), yet, we make sure not to use more than readily-available consumer equipment, like a camera with flash. Starting from a SfM-generated mesh, we apply PS to refine both geometry and reflectance, where the latter is expressed in terms of data-driven Bidirectional Reflectance Distribution Function (BRDF) representations. We also introduce a novel approach to infer complete BRDFs starting from the sparsely sampled data-driven reflectance information captured with this setup. Our approach is experimentally validated by modeling several challenging objects, both synthetic and real.*

## 1. Introduction

Producing photo-realistic renderings of objects calls for more than high-quality 3D shape acquisition. Obtaining a detailed model of reflectance across the surface is equally crucial. Yet, it is fair to say that 3D shape extraction has progressed far more than the extraction of reflectance. Nowadays, un-calibrated SfM comes close to what laser scanners can achieve. Furthermore, several PS-based methods have been proposed that take advantage of material characteristics to improve results where traditional SfM might fail. In contrast, systems that provide a reflectance model close to the ideal BRDF at the different surface points are by and large lacking still. Most systems go for a far simpler reflectance model, like Lambertian behavior, possibly mixed with some specular lobes. Therefore, 3D shape and surface reflectance capture remains an important and challenging problem. Typically, this capture is performed with sophisticated hardware setups which have proven to be accurate, but complicated and inaccessible to casual users.

The primary contributions of this paper are summarized as follows: (1) Our setup is restricted to the use of a flash-equipped camera, providing to casual users the abil-



Figure 1. Top: s-o-a un-calibrated 3D reconstruction, with the shape rather smooth and Lambertian reflectance rendering. Bottom: proposed method with its refined geometry and extracted BRDF map.

ity to extract both 3D shape and surface reflectance from readily-available consumer equipment. The user should walk around the object, taking photos under the illumination of the camera’s flash. Using un-calibrated SfM (and dense matching) an initial 3D mesh is extracted. (2) The method then estimates realistic lower-dimensional BRDFs (BRDF slices) from sampled sections<sup>1</sup>, and iteratively uses the estimates to refine both reflectance and geometry based on a new optimization technique. (3) Finally, we present an algorithm to infer complete BRDFs from the sampled sections. Fig. 1 gives a preview.

## 2. Previous work

The amount of research on 3D shape and reflectance acquisition is vast. We organize the overview by topic:

**3D shape acquisition.** PS computes surface normals from shading variations observed in images taken under differently-oriented illumination. Recent efforts have been

<sup>1</sup>Since the amount of viewpoints/lights is limited, the BRDF is only sparsely sampled (see Sec. 4)

directed towards robustification against outliers, such as shadows and specular highlights [5, 27]. They exploit the fact that most specular materials exhibit an approximately Lambertian behaviour for at least a subset of the viewing/lighting combinations, to single out this matte component and apply Lambertian PS. Other methods [14, 15] are dedicated to capturing non-Lambertian phenomena such as specular highlights and iridescence. In general, the resulting fixed texture map acquired with all these methods is insufficient to represent reflectance properly.

**Flash-aided reconstruction.** Flash-aided reconstruction offers a number of advantages, like data capture under low ambient light, minimum number of self-shadows and reasonably-controlled illumination. Several works [28, 2] fuse flash/no-flash image pairs to remove undesired artifacts like noise and specular highlights. Melendez et al. [24] use shading information from flash images for Lambertian PS and Lanman et al. [19] exploit specular highlights via a calibrated multi-flash system to obtain shape and reflectance using SfM and PS. Our data-driven approach makes far weaker assumptions both on geometry and reflectance and does not require any calibration target in the scene.

**Reflectance acquisition.** BRDFs describe the fraction of reflected light for all pairs of incoming/outgoing light directions and are traditionally acquired with dedicated hardware setups, e.g. [23], assuming known shapes. Methods like [10, 30] are related to ours but are only applicable to near-planar surfaces.

**3D shape scanning and reflectance estimation.** Laser scanners and Structured-Light (SL) patterns can be used to obtain accurate 3D shape [32, 9]. Parametric reflectance models can then be fitted to each surface point based on image observations [20, 21], but solving the difficult problem of precisely registering images with 3D shape is crucial otherwise artifacts appear in misaligned regions. Combining reflectance from PS and shape from SL [39, 25, 3] solves this problem but the complicated setups limit the method’s applicability to casual users.

**Joint 3D shape and reflectance estimation.** Because of the complementarity between shape and reflectance, quite some research effort has been put into jointly estimating shape and reflectance from imagery, or take advantage of the reflectance characteristics to improve the SfM result, using principles from PS. Methods like [13, 12, 26] fit specific parametric BRDF models to input data, which may result in performance degradation when objects have a reflectance different from the assumed model. Some other methods achieve highly accurate results by employing sophisticated hardware. Ma et al. [22] and Ghosh et al. [11] used a light stage where LEDs’ intensity was precisely controlled, whereas Holroyd et al. [18] required coaxial lights. Both setups need expensive and complicated hardware to work. Some recent algorithms exploit various reflectance charac-

teristics, like symmetries [4, 17] or monotonicity [34], to estimate shape and reflectance, but they either require up to a thousand input images [17] or rely on fragile optimization [4, 34]. Tan et al. [35] and Chandraker et al. [6] recovered iso-contours of depth and gradient magnitude for isotropic surfaces having the disadvantage that they need additional user interactions or boundary conditions to recover the shape. Zhou et al. [41] have proven that the iso-contour approach for PS can be properly integrated with SfM, using a simple (yet fully-calibrated) lighting setup. Higo et al. [16] used a hand-held camera with a single light source to recover shape and reflectance from images but their approach relies on the Lambertian part of the material alone, and focuses on shape rather than material.

### 3. System overview

We investigate a simple setup consisting of a camera with built-in flash. The flash point of view is typically very close to the camera point of view. Compared to traditional PS approaches [41, 30] however, we will not have any reflectance information from light sources oblique to the camera point of view. Also, having both light and camera viewpoint aligned limits to undo the scale ambiguity in SfM [16]. If only one single light source is available, other assumptions are needed to recover reflectance information. We adopt the observation that many objects are built out of a limited set of materials [20], meaning that one can still observe a material’s reflectance changes by observing all points of the same material in the different camera viewpoints.

Given an image sequence, the system is initialized using un-calibrated SfM to recover the extrinsic and intrinsic camera parameters and reconstruct an initial 3D point-cloud or mesh (see Sec. 4). While many methods combining SfM and PS, use isotropy and symmetry assumptions on the observed BRDFs [6, 4], or limit the computation to the Lambertian part of the observed reflectance [41, 16], in this paper we investigate a data-driven approach, going straight for a proper BRDF estimation without any parametric or other assumptions, by observing the reflectance in the camera viewpoints, including specular highlights. The latter prove to show intricate surface deviations that can not be observed by angles obeying the Lambertian law (e.g. [40, 8]). Here we want to remind that the overall goal is not only to arrive at an improved geometry and normal distribution, but also a plausible BRDF representation that can be used to create realistic renderings of the resulting 3D model.

Every point in the mesh is characterized by its position, its normal, and a set of weights related to a limited set of base materials, each with a BRDF representation. In this paper we introduce a method where each of these parameters is to be refined/optimized by minimizing the color difference between the original and rendered images, using the BRDF estimations. Obviously, for our flash-based setup,

we can not recover the full BRDF space, but we can recover a lower-dimensional BRDF (as will be discussed later). In order to arrive at a better approximation of the full BRDF, we propose a method to infer the missing 2D BRDF information from the 1D slice, allowing the final model to be rendered more confidently when the lighting conditions are different from the camera viewpoint.

#### 4. Initial geometry and reflectance model

**Input assumptions and geometry.** The input to our method is a sequence of  $M$  images. Using SfM, both an initial shape as well as the intrinsic and extrinsic camera parameters are recovered. There are no stringent requirements on the types of materials that can be processed, but generally for the approach to work, they need to show a diffuse component and/or a reasonable amount of texture variation. Nevertheless, we tackle quite challenging examples showing a lot of reflectivity as well as texture-less smooth parts. Pure mirror surfaces have to be ruled out. It is useful to note that the flash is treated as a point light source and that we assume it to be dominant over other illumination; possible effects of environment lighting have not been considered. It is also assumed that the response of the camera is linear, a characteristic that holds for the RAW images captured with the cameras used in our experiments.




Figure 2. A preview on the whole pipeline, from left to right: Ground truth image, SfM-based 3D features, initial mesh, final refined mesh, and final rendering.

Our particular implementation of the SfM scheme [37], involves the detection of SURF-based features in each image, which are matched throughout the sequence. The feature correspondences determine the intrinsic parameters, pose estimations and 3D positions of these features. In a subsequent step, image pairs showing enough correspondences are rectified and disparity search is performed based on dynamic programming to create dense depth-maps. Using belief propagation on the depths observed in the original images, an initial dense mesh can be constructed. A preview of the whole pipeline, including the SfM step is shown in Fig. 2. It should be noted that for the following experiments, alternative approaches have been verified. Using the 3D positions of the SURF features alone, an initial mesh can be constructed by simple Poisson reconstruction. The

latter shows lower quality, but serves the purpose of verifying the sensitivity of the optimization process with respect to the initial mesh quality. The same goes for the chosen mesh resolution after dense matching.

For the practical implementation of the consecutive BRDF estimation, the initial mesh is further subdivided (if needed) to match the resolution of the images as seen from the camera viewpoint. Thus, the vertices of the mesh serve as a point-cloud for which both the normals, positions and BRDF curves are to be refined.




Figure 3. Left: material sample and corresponding 2D BRDF and BRDF slice. Right: initial BRDF slices obtained by clustering the input observations.

**Lower-dimensional BRDFs.** Our starting point is the BRDF representation introduced by Romeiro et al. [31]. They reduce the 4D BRDF to 2D, because it can effectively capture the behavior of the majority of real materials. This parameterization is based on half and difference angles  $\theta_h$  and  $\theta_d$ . The half vector  $h$  is the bisector of lighting  $\omega_i$  and viewing  $\omega_o$  directions, thus  $\theta_h$  measures the angular deviation from the direction of ideal specular reflection. This can be seen in the 2D BRDF of Fig. 3 (left), where the top row corresponds to  $\theta_h = 0$ . The abscissa represents  $\theta_d$ .

With  $B$  as our per color-channel 2D BRDF, we model the reflected RGB color vector as

$$\tilde{c} = \max(0, n \cdot h) B(\theta_d, \theta_h) \quad (1)$$

where the clipped dot product accounts for attenuation due to an off-normal incident light direction (Lambert’s cosine law). For our fixed camera-flash setup there is also the observation that  $\omega_i \approx \omega_o$ , i.e.  $\theta_d \approx 0$ . The sole remaining parameter  $\theta_h$  indicates the direction of both camera and flash relative to the local surface normal. In Fig. 3 (left) our samples populate the first column of the 2D BRDF (denoted by blue). This is not just a column of the 2D BRDF, it contains vital information about the dynamic range. Indeed, for most materials the maximal intensity is obtained when the viewing, lighting, and surface normal directions are aligned. This column will be referred as 1D BRDF or BRDF slice.

**Base materials.** Assuming a material can be observed for many points (with a variety of orientations) in the 3D model, e.g. the sphere in Fig. 3 (left), a single image can provide many samples of the 1D BRDF. To represent

the BRDF, analytic models have been popular in literature [26], but usually involve the non-trivial guesstimation of initial parameters, while being susceptible to noise. We start from a purely data-driven approximation of the first column  $B(\theta_h)$  ( $= B(0, \theta_h)$ ), expressed as a  $\theta_h$ -indexed vector of RGB values. Yet, we have to make sure not to mix up samples observed at points with different material properties. Therefore, we follow an approach similar to the one introduced by Lensch et al. [20]. We assume our surface to be composed of a finite number of base materials and let all surface points belonging to the same material contribute to the first column of that base material’s BRDF.

**Clustering.** Consider an initial 3D model consisting of  $N$  3D points  $P$  and normals  $n$ . We project every  $P_i$  into each of the camera viewpoints recovered from SfM, to obtain  $M$  pairs of RGB measurements ( $c$ ) and half angles ( $\theta_h$ ):

$$\forall \{P_i, n_i\} \mid_{i=1 \dots N} : C_i = \{c, \theta_h\}_{ij} \mid_{j=1 \dots M, \zeta(i,j)=1}. \quad (2)$$

The visibility  $\zeta_{ij}$  equates to one if  $P_i$  is visible in image  $j$ , and to zero otherwise. From each  $C_i$ , we compute a low quality 1D BRDF, by quantizing the half angle space into  $Q'$  equally sized bins over  $\theta_h = [0 \dots \pi/2]$  and taking a weighted sum of all color observations that belong to the same bin. In our experiments,  $Q' = 3$ . We use the method proposed by Tingdahl et al. [36] to partition these low quality 1D BRDFs into  $K$  base materials. To determine the total number of materials, we run weighted k-means with an increasing number of base materials  $k$ , and stop when the newly generated base BRDFs are too close to the ones already created in the previous step (see [36]). For the rooster in Fig. 3 (right), this approach results in 3 base materials.

For each base material  $k$ , we now re-sample the BRDF slice using all of its assigned measurement pairs:  $\{c, \theta_h\}_{ij} \mid i \in k$ . Again, we quantize the half angle space into  $Q$  equally sized bins (in our experiments  $Q = 30$ ). Fig. 3 (right) shows the 3 initial base BRDFs for the rooster.

## 5. Reflectance and geometry refinement

Due to lack of (initial) structure, (specular) reflections and noise, there are inaccuracies on the initial 3D points and normals and the initial 1D BRDFs may underestimate possible reflections. In a next step, we would like to refine the point positions, local normals, and (weighted) BRDF representations, such that the estimated reflectance for each point in each image fits the observations. Obviously, these are a lot of parameters to take into account. In order to control this process, we propose an optimization method that is alternating between optimizing 1D BRDFs, normals, the weighting between materials, and finally points positions.

Given an initialization for 3D points, normals, weights and 1D BRDFs from the previous sections we aim to adjust the state  $x = \{B_1, \dots, B_k, n_1, \dots, n_N, w_{11}, \dots, w_{kN}, p_1, \dots, p_N\}$ .

To this end, we minimize a cost function:

$$\sum_{i=1}^N \sum_{j=1}^M \zeta(i, j) d(\tilde{c}_{ij}(x), c_{ij}) + \lambda_B E_B + \lambda_n E_n + \lambda_w E_w \quad (3)$$

with  $d$ , the difference between modeled  $\tilde{c}_{ij}$  (from Eq. 1) and observed measurements  $c_{ij}$  and  $\zeta(i, j)$  the visibility function. Below we describe the different constraints that are taken into account:

**1D BRDF.** The two BRDF constraints  $E_B = \lambda_1 E_1 + \lambda_2 E_2$  will be explained shortly after describing why such a BRDF curve refinement is crucial. The true BRDF curve might very well extend beyond the saturation point imposed by the image measurements (which is 1 in our case), especially when close to a specular highlight ( $\theta_h \approx 0$ ). Thus, saturated image measurements only provide limited information regarding the true shape of the curve. For all we know, the true BRDF values can be anywhere above 1. Therefore  $d$  may prevent the BRDF to exceed 1 if we attempt to fit these measurements directly. At the other end of the curve, the quality of the image measurements deteriorates as the view-to-surface angle gets about parallel to the surface. Therefore, we discard any measurements for which  $\cos \theta_h < 0.15$ . In summary, our BRDF curve has two loose ends that we need to somehow tie them up.

*Smoothness ( $E_1$ ):* Working in  $L^*a^*b^*$  space, we encourage a low second derivative of the luminance component  $L$ , and a low first derivative of the chromaticity components  $a^*$  and  $b^*$ . This reflects the fact that most commonly encountered BRDFs show only a minor change in color, but may exhibit large variations in intensity.

$$E_1 = \frac{1}{Q} \sum_{q=1}^Q (\nabla^2 B_L^2(q) + \nabla B_{a^*b^*}^2(q)). \quad (4)$$

One might argue that apart from the surface albedo, the reflective color and incident light color might be quite different, however these constraints control erratic introduction of false colors nearby the specular peaks, which can be very disturbing in the final renderings.

*PCA reconstruction error ( $E_2$ ):* While the smoothness constraint has a positive effect on the overall shape of the curve, the behavior might still be undesired for regions with low or no measurements. For instance, the curves can be arbitrarily extended for any  $\cos \theta_h < 0.15$  since no valid measurements exists beyond that point. We predict the behavior in such regions by encouraging similarity to BRDFs sampled in the real world. This is expressed in terms of an error between the observed curve  $B$ , and a reconstructed curve  $\tilde{B}$ , obtained by projecting  $B$  into a basis calculated by applying PCA to the 1D BRDF samples of MERL database [23]. We used 5 principal components.

$$E_2 = \frac{1}{Q} \sum_{q=1}^Q (\tilde{B}(q) - B(q))^2. \quad (5)$$




Figure 4. In the refined BRDF (right) the specular highlight is more prominent than in the initial (left), which gives a visual improvement in the final renderings of the rooster of Fig.2

**Normal smoothness.** We encourage the local surface curvature to be constant. This is approximated as the difference of the angles between the normal of each vertex and those of the neighboring vertices.

$$E_n = \frac{1}{V} \sum_{v=neighbors(i)} (\arccos(n_i n_v))^2. \quad (6)$$

**BRDF weighting.** The individual  $B$  of each point is expressed as a weighted sum of  $K$  BRDF slices  $B_k$ . We want to make sure that the current weights estimates  $\tilde{w}$  will not deviate too much from their initially assigned clusters  $w$ .

$$E_w = \frac{1}{K} \sum_{k=1}^K (\tilde{w}(k) - w(k))^2 \quad (7)$$

**3D points positions.** Having adjusted the normals, the new estimates can be used to improve the positions of the object's points. For this particular task we rely on the method proposed by [25]. To this end we minimize

$$E_p = (1 - \lambda_p) \sum_{i=1}^N \sum_{j,k \in \Omega_i} [n_i^m \cdot (P_j - P_k)]^2 + \lambda_p \sum_i [T_i \cdot (P_i - P_i^0)]^2 \quad (8)$$

The first term allows points to move such that, the tangent plane formed by the neighbors  $\Omega_i$  of each point  $i$ , shows a normal close to the estimated normal. Alternatively stated, we favor neighboring edges which are orthogonal to the current normal estimations. Additionally, in order to avoid self-intersections, the second term favors motion in the normal direction, and penalizes motion in the tangent plane.  $T_i = \phi n_i^m n_i^{mT} + \xi(I - n_i^m n_i^{mT})$  (see [25]).

**Minimization details.** We minimize the different constraints using non-linear least squares minimization [1]. As indicated, since a full global refinement is computationally expensive we optimize in 3 steps. First, we refine  $B_k$  and  $n_i$  where  $i \in k$  for each cluster individually keeping  $w_i$  constant ( $\lambda_w = 0$ ). Fig. 4 shows how the initial BRDF observation of the rooster has converged to a state with a high specularity and MERL-like tail. In a second step we fix  $B$ ,




Figure 5. 1D-2D BRDF prediction pipeline.

and refine  $w_i$  and  $n_i$  for each individual point  $i$  ( $\lambda_B = 0$ ). Finally we update the 3D positions using  $E_p$ .

## 6. 2D BRDF inference

In the former sections we estimated 1D BRDFs. The question is whether one can infer reasonable approximation for the missing  $\theta_d$  of the 2D BRDF. In literature,  $\theta_h$  is accounted for by the microfacet normal distributions (NDF) and  $\theta_d$  is basically related to the Fresnel effect. Having no clue on the dependency of  $\theta_d$  for this setup, a common assumption in literature is that the reflection characteristics of the object remain largely stable with  $\theta_d$  (e.g. [7]). Looking at  $(\theta_d, \theta_h)$  space of real-life material samples such as MERL, we see a kind of common behavior that seems statistically relevant. In this paper we go beyond the simple approximation of assuming that the reflection characteristics remain stable over  $\theta_d$  [7] and we try to predict the dependence over  $\theta_d$ .

To determine possible statistical correlation we first perform PCA on both the 1D BRDF slices ( $\theta_h$ ) and the corresponding 2D BRDFs ( $\theta_d, \theta_h$ ) of all MERL samples. We work on log space (natural logarithm), in order to deal with the large differences between the specular and non-specular values of the measured BRDFs [23]. To find a common embedding manifold where the correlation between the PCA representations of the 1D slices and 2D BRDFs is maximized, we apply Canonical Correlation Analysis (CCA) to the parameters derived from the two PCAs. We found a statistical correlation of [0.8, 0.9] across the principal components, which indicates the relevance of such a prediction. Given the two new bases from CCA we can project the PCA parameters of the 1D columns and 2D BRDFs into this common CCA space. Gaussian Process Regression (GPR) [29] of the projected parameters in the common CCA space results in a prediction matrix, converting 1D to 2D BRDFs. Fig.5 gives the block diagram of our approach.

We chose to cluster the MERL BRDFs into groups of similar statistical behavior. Strikingly, using an iterative approach based on Spectral Clustering [38] we automatically found 2 clusters, which represent the 'lambertian' and 'specular' materials, resp. In these clusters statistical correlation between 1D and 2D BRDFs is maximal. We end up

with a set of prediction matrices, each for the different clusters. For a new sample, the proper cluster is the one with the smallest PCA reconstruction error. We have observed that for this cluster the prediction is indeed optimal (Fig 6 (J)).

The prediction step has important consequences for rendering purposes. As many of the original MERL samples show, there are quite some highlights in the higher  $\theta_d$  range, which can not be represented by a single 1D BRDF slice copied over the remaining reflectance space [7]. These areas come into play especially when rendering the model in environment lighting, as illustrated in Fig 6 (I). We validate our approach using 10-fold cross-validation. In Fig. 6 (K), we show the error difference between the ground truth MERL BRDFs, and resp. the estimate assuming constancy over  $\theta_d$  (green) [7], and our predicted profiles (blue for specular, purple for lambertian materials). The testing samples are always excluded from the training set. Overall, the error decreases substantially. As a final note: if the material shows color changing effects over  $\theta_d$ , the prediction is not able to recover that (see last row of Fig 6 (I)).

## 7. Results

The pipeline has been tested on various, synthetic and real, challenging examples. Unless stated otherwise, we used 100 images of 12M resolution for each experiment. All the rendered images are taken from viewpoints that have not been used in the optimization pipeline. All figures are composited in Fig. 6, (A)-(N).

A synthetic sphere is shown in (A) with a reflective MERL material (MERL id 68). Noise has been added on the normals up to 45 degrees from their ground truth orientation, and on the positions up to 10% of the radius. For these noise levels, the optimization scheme was able to recover the shape within 0.3 RMSE. (B) shows how the initial BRDF estimate is seriously affected by the noise, but recovered after refinement. The specular peak is slightly lower, but sufficient to provide a proper render. The same sphere was covered with two MERL materials (MERL id 68 and 98) in (C), and went through the full pipeline including SfM. In (D), additional concavities and convexities were added without texture, so only photometric input can play a role. In both cases the recovered 3D shape was within 0.5 RMSE of the radius. Artifacts are mainly observed on the cluster boundaries and the estimated BRDFs (E) coincide well with the original MERL profiles up to a scale factor.

Areas (F), (G), (H) show some experiments on real examples. The Pericles figurine (F) is a bronze-like statue recorded with a Canon 600D. It has a distinctive specular reflection, and also shows inter-reflections. The top row illustrates how was the overall shape improved and compares the final renderings with the ground truth. For this particular object we also tested incoming light orthogonal to the viewing direction (L). Such a lightning condition is very

different from the head-on (flash-based) lighting setup. So here we rely on the 1D-2D prediction to approach realistic visualization. The side light renderings show that the predicted 2D BRDF is definitely improving with respect to the initial 1D slice. The environment renderings (N) highlight the substantial difference when the object is rendered with the 1D BRDF (slice), instead of the 2D prediction.

The Buddha statue (G), originating from the OBJECTS2011 data set [33], was captured with a multi-camera, multi-light dome. We mimicked our fixed-flash scenario, singling out the images for which the flash and camera were aligned, resulting in a subset of 107 images. The object shows inter-reflections, changing chromaticities and mixing materials which make the material boundaries less clear. The results on shape and reflectance are nevertheless convincing. An environment render is shown in (N).

The bronze rabbit (H) is recorded using an Android tablet, with built-in flash. The surface shows typical cracks and irregularities of the bronze material. Although improved, some errors are located at the back of the ears and the top of the back, due to occlusions.

The examples show that we are able to estimate reliable 3D shape information and BRDF slices, and create virtual renderings that match the observations, despite of the complexity of the overall pipeline. The method relies on a careful optimization between the observations and the estimated BRDFs, local normal characteristics and geometry. Below we list a few observations on the processing aspects. The constraints parameters range over  $\lambda_1 \approx [0.1, 1.0]$ ,  $\lambda_2 \approx [0.1, 0.5]$ ,  $\lambda_n \approx [0.1, 0.5]$ ,  $\lambda_w \approx 0.1$ ,  $\lambda_p \approx [0.1, 0.3]$ ,  $\gamma \approx 0.1$ . For real examples showing a higher complexity in texture and reflectance, resulting in higher noise levels in the initial observations, higher parameter values are recommended. With respect to clustering, we have observed that the quality of the results is not only dependent on the amount of clusters but also on the amount of observations per cluster. Finally, on the level of materials statistics, which we used both in the optimization process and the 2D BRDF prediction, we have only included a limited amount of samples. The 100 (MERL) material samples are probably not representative to cover all possible materials and their variability. On the computational side, the processing time depends linearly on the amount of 3D points, the image resolution and the amount of images.

## 8. Conclusion

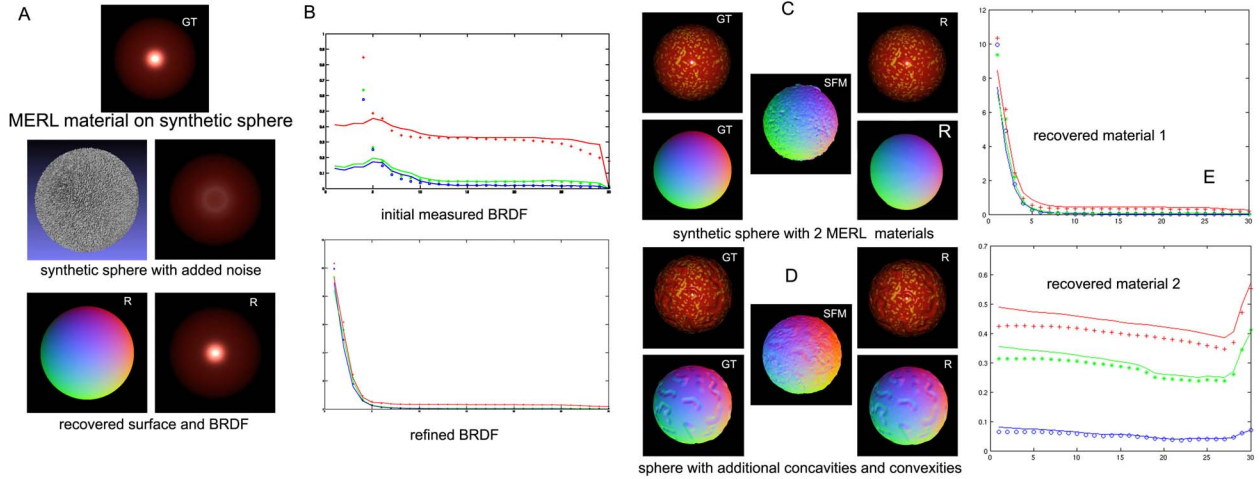
We presented a method to determine surfaces and material characteristics of objects using a camera with flash as the sole acquisition device. We have obtained promising results for synthetic and real objects, which encourages us to invest in a more detailed analysis of the work-flow, material statistics, alternative methodologies for the optimization and possibly alternative lighting configurations.

**Acknowledgements.** This research is supported by ROV-INA project which is funded from the EU Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 600890.


## References

- [1] S. Agarwal, K. Mierle, and Others. Ceres solver. <http://ceres-solver.org>. 5
- [2] A. Agrawal, R. Raskar, S. K. Nayar, and Y. Li. Removing photography artifacts using gradient projection and flash-exposure sampling. *SIGGRAPH*, 2005. 2
- [3] D. G. Aliaga and Y. Xu. Photogeometric structured light: A self-calibrating and multi-viewpoint framework for accurate 3d modeling. *CVPR*, 2008. 2
- [4] N. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. *CVPR*, 2008. 2
- [5] M. Chandraker, S. Agarwal, and D. Kriegman. Shadowcuts: Photometric stereo with shadows. *CVPR*, 2007. 2
- [6] M. Chandraker, J. Bai, and R. Ramamoorthi. A theory of differential photometric stereo for unknown isotropic brdfs. *CVPR*, 2011. 2
- [7] M. Chandraker and R. Ramamoorthi. What an image reveals about material reflectance. *ICCV*, 2011. 5, 6
- [8] T. Chen, M. Goesele, and H. P. Seidel. Mesostructure from specularity. *CVPR*, 2006. 2
- [9] J. Davis, D. Nehab, R. Ramamoorthi, and S. Rusinkiewicz. Spacetime stereo: A unifying framework for depth from triangulation. *PAMI*, 2005. 2
- [10] Y. Dong, J. Wang, X. Tong, J. Snyder, Y. Lan, M. Ben-Ezra, and B. Guo. Manifold bootstrapping for svbrdf capture. *SIGGRAPH*, 2010. 2
- [11] A. Ghosh, T. Chen, P. Peers, C. A. Wilson, and P. Debevec. Estimating specular roughness and anisotropy from second order spherical gradient illumination. *EGSR*, 2009. 2
- [12] D. B. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and spatially-varying brdfs from photometric stereo. *PAMI*, 2010. 2
- [13] C. Hernandez, G. Vogiatzis, and R. Cipolla. Multi-view photometric stereo. *PAMI*, 2007. 2
- [14] A. Hertzmann and S. M. Seitz. Example-based photometric stereo: shape reconstruction with general, varying brdfs. *PAMI*, 2005. 2
- [15] T. Higo, Y. Matsushita, and K. Ikeuchi. Consensus photometric stereo. *CVPR*, 2010. 2
- [16] T. Higo, Y. Matsushita, N. Joshi, and K. Ikeuchi. A handheld photometric stereo camera for 3d modeling. *ICCV*, 2009. 2
- [17] M. Holroyd, J. Lawrence, G. Humphreys, and T. Zickler. A photometric approach for estimating normals and tangents. *SIGGRAPH*, 2008. 2
- [18] M. Holroyd, J. Lawrence, and T. Zickler. A coaxial optical scanner for synchronous acquisition of 3d geometry and surface reflectance. *SIGGRAPH*, 2010. 2
- [19] D. Lanman, P. G. Sibley, D. Crispell, Y. Zhao, and G. Taubin. Multi-flash 3d photography: capturing shape and appearance. *SIGGRAPH*, 2006. 2
- [20] H. P. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H. P. Seidel. Image-based reconstruction of spatial appearance and geometric detail. *TOG*, 2003. 2, 4
- [21] S. Lombardi and K. Nishino. Single image multimaterial estimation. *CVPR*, 2012. 2
- [22] W. C. Ma, T. Hawkins, P. Peers, C. F. Chabert, M. Weiss, and P. Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. *EGSR*, 2007. 2
- [23] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *SIGGRAPH*, 2003. 2, 4, 5
- [24] F. Melendez, M. Glencross, G. J. Ward, and R. J. Hubbard. High-resolution relightable buildings from photographs. *SIGGRAPH*, 2011. 2
- [25] D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi. Efficiently combining positions and normals for precise 3d geometry. *SIGGRAPH*, 2005. 2, 5
- [26] G. Oxholm and K. Nishino. Shape and reflectance from natural illumination. *ECCV*, 2012. 2, 4
- [27] J. Park, S. N. Sinha, Y. Matsushita, Y. W. Tai, and I. S. Kweon. Multiview photometric stereo using planar mesh parameterization. *ICCV*, 2013. 2
- [28] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama. Digital photography with flash and no-flash image pairs. *TOG*, 2004. 2
- [29] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2005. 5
- [30] P. Ren, J. Wang, J. Snyder, X. Tong, and B. Guo. Pocket reflectometry. *SIGGRAPH*, 2011. 2
- [31] F. Romeiro, Y. Vasilyev, and T. Zickler. Passive reflectometry. *ECCV*, 2008. 3
- [32] S. Rusinkiewicz, O. Hall-Holt, and M. Levoy. Real-time 3d model acquisition. *SIGGRAPH*, 2002. 2
- [33] C. Schwartz, M. Weinmann, R. Ruiters, and R. Klein. Integrated high-quality acquisition of geometry and appearance for cultural heritage. *VAST*, 2011. 6
- [34] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Elevation angle from reflectance monotonicity: Photometric stereo for general isotropic reflectances. *ECCV*, 2012. 2
- [35] P. Tan, L. Quan, and T. Zickler. The geometry of reflectance symmetries. *PAMI*, 2011. 2
- [36] D. Tingdahl, C. Godau, and L. V. Gool. Base materials for photometric stereo. *ECCV*, 2012. 4
- [37] D. Tingdahl and L. V. Gool. A public system for image based 3d model generation. *MIRAGE*, 2011. 3
- [38] R. Vidal. Subspace clustering. *IEEE Sign. Proc. Mag.*, 2011. 5
- [39] L. Zhang, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and motion under varying illumination: Unifying structure from motion, photometric stereo, and multi-view stereo. *ICCV*, 2003. 2
- [40] Z. Zheng, L. Ma, Z. Li, and Z. Chen. An extended photometric stereo algorithm for recovering specular object shape and its reflectance properties. *Comp. Sc. and Inf. Sys.*, 2010. 2
- [41] Z. Zhou, Z. Wu, and P. Tan. Multi-view photometric stereo with spatially varying isotropic materials. *CVPR*, 2013. 2

Synthetic Examples



Real Examples



2D BRDF Inference

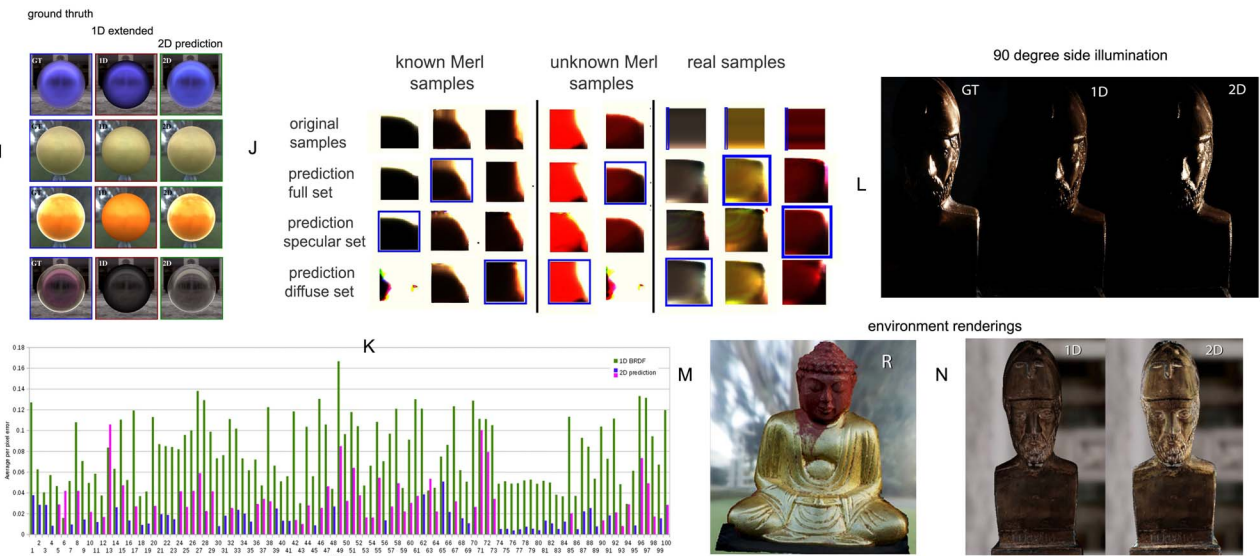


Figure 6. Synthetic and real experiments on recovering shape, reflectance and BRDF prediction. Note the initials on top of the images: GT = ground truth, R = result, SFM = structure from motion, C = cluster