

On-Site Example-Based Material Appearance Acquisition

Y. Lin¹ P. Peers² A. Ghosh¹

¹Imperial College London
²College of William & Mary




Figure 1: Example-based material appearance estimation of a blue faux leather purse in an outdoor environment. (a) Input HDR photograph of the exemplar. (b) Corresponding light probe. (c) A sphere rendered under the same illumination with the estimated BRDF. (d) Synthesized master tile of the material's mesostructure: top - surface normals, bottom - specular reflection occlusion. (e) Rendering with both the estimated BRDF plus synthesized mesostructure details.

Abstract

We present a novel example-based material appearance modeling method suitable for rapid digital content creation. Our method only requires a single HDR photograph of a homogeneous isotropic dielectric exemplar object under known natural illumination. While conventional methods for appearance modeling require prior knowledge on the object shape, our method does not, nor does it recover the shape explicitly, greatly simplifying on-site appearance acquisition to a lightweight photography process suited for non-expert users. As our central contribution, we propose a shape-agnostic BRDF estimation procedure based on binary RGB profile matching. We also model the appearance of materials exhibiting a regular or stationary texture-like appearance, by synthesizing appropriate mesostructure from the same input HDR photograph and a mesostructure exemplar with (roughly) similar features. We believe our lightweight method for on-site shape-agnostic appearance acquisition presents a suitable alternative for a variety of applications that require plausible “rapid-appearance-modeling”.

CCS Concepts

• Computing methodologies → Reflectance modeling;

1. Introduction

In realistic digital content creation pipelines, artists often wish to closely match the appearance of virtual objects to the appearance of real world materials and objects. Such a content creation process typically requires capturing physical reflectance properties of the desired materials. However, traditional reflectance acquisition approaches are expensive, time consuming, and require dedicated setups, hindering adoption by non-experts. Recently, a number of convenient and accessible appearance modeling methods [AWL15, AAL16, RPG16, HSL*17] have been proposed that use a mobile phone for on-site acquisition of the appearance of

spatially varying materials. However, these methods impose strong restrictions on the sample's shape (e.g., planar), restrict the spatial variation (e.g., texture-like), or require a significant amount of calibrated data.

In this paper, we propose a novel on-site example-based material appearance acquisition technique that only requires a single HDR photograph of an object with a homogeneous material exemplar under *known* but uncontrolled natural lighting. We do not require the exemplar to be planar or to have a known shape (other than requiring it to be convex, or exhibiting sufficient shape variation due to surface mesostructure), nor does our method explicitly recover the

unknown shape, greatly simplifying the acquisition process for a non-expert user. Our method only makes modest assumptions on the underlying material and lighting: we assume the material is dielectric and can be described by an isotropic BRDF, with a possible stationary texture-like mesostructure, and that the incident lighting is uncontrolled, but dominated by a single bright region with colorful surrounding ambient illumination. Given an input HDR photograph of the exemplar object (next to a color chart for calibration) and a light probe of the incident illumination, we estimate parameters of a microfacet BRDF model [CT82] that best describes the homogeneous material appearance. Our estimation method is agnostic to the underlying surface mesostructure, enabling us to estimate the BRDF from both smooth surfaces as well as for objects with a texture-like mesostructure. While such a BRDF estimate is sufficient for reproducing the appearance of objects with a smooth surface, for objects with a texture-like appearance we additionally also synthesize a repeating surface mesostructure. We quantitatively analyze the robustness of our method on synthetic scenes with respect to variations in object shape, mesostructure, lighting, and BRDF parameters, and qualitatively demonstrate on a variety of real-world examples that the estimated BRDF in conjunction with synthesized mesostructure (where appropriate) can faithfully reproduce appearance for a variety of materials (e.g., see Figure 1).

2. Related Work

We focus this overview on related work that shares some of the goals of our method: in-situ appearance acquisition, appearance modeling under natural lighting with unknown shape, recovering meso-structure, and exploiting chromatic cues for scene modeling. For a detailed overview of appearance modeling we refer to [DRS08, WK15].

In-situ Acquisition In-situ acquisition and estimation of appearance has received significant research attention in the past few years. A majority of methods require multiple photographs of the material sample either under active lighting [AWL15, RPG16, HSL*17, XNY*16] or aided by prior knowledge on the shape of the scene [DCP*14, PCDS12]. In this paper, we pursue a method that only requires a single photograph of the scene under uncontrolled natural lighting with unknown shape.

The methods of Romeiro et al. [RVZ08, RZ10], Lombardi and Nishino [LN16], and Zhou et al. [ZCD*16] estimates reflectance properties from an object with known shape under either known or unknown natural lighting. We observe that in practice, capturing the incident lighting is easier than capturing and aligning the shape. We therefore assume prior knowledge of the lighting, but not of the shape.

A number of methods have been proposed that rely on deep learning to estimate the reflectance properties and meso-structure from a single photograph under flash lighting [DAD*18, LSC18, AAL16, LXR*18] or under natural lighting [LDPT17, YLD*18]. While promising, these methods either control the incident lighting, and/or are limited to planar samples only.

Appearance Modeling under Natural Lighting with Unknown Shape Estimating both shape and reflectance under natural illumination is an ill-conditioned problem. Existing methods either re-

quire multiple observations of the scene [BM15, XDPT16], possibly in conjunction with scene depth [WWZ16], or assume a smooth surface [ON16, GRR*17, MMZ*18]. However, jointly estimating shape and reflectance suffers from the potential pitfall that inaccuracies in shape estimation affect the accuracy of the reflectance estimation. Our method takes a different approach and foregoes estimating the shape all together. Instead we estimate the reflectance properties by exploiting statistical properties of the materials.

Recovering Meso-Structure Photometric stereo [Woo80] is a popular method for estimating the meso-structure in the form of a normal map from just a few photographs under different directional lighting conditions. Subsequent research has proposed numerous extensions, including normal estimation under natural lighting [BJK07]. However, classic photometric stereo is limited to diffuse materials only. Example-based methods [HS03, THS04] overcome this limitation by also capturing an exemplar of the same materials and with known geometry, or simply matching specular highlight peaks on the material due to a point light source against those seen on a sphere [CGS06]. These method only focus on meso-structure estimation. In contrast, our solution also recovers surface reflectance and does not require controlled point source illumination.

Wang et al. [WSM11] use step-edge lighting to estimate the surface reflectance of a homogeneous material as well as the meso-structure. While they pursue a similar goal as us, their method is limited to planar surfaces and requires active illumination. In contrast, we do not require active lighting and can resolve the reflectance of material samples with arbitrary shapes.

Aittala et al. [AWL15] employ a combination of observations under flash illumination and ambient lighting to recover spatially varying surface reflectance properties of planar samples with a texture-like appearance. In this work, we borrow similar mesostructure exemplars from their measured database for synthesizing appropriate mesostructure on a target material with a texture-like appearance.

Chromatic Cues for Scene Modeling Our method exploits chromatic variations in the lighting to infer reflectance properties. Chromatic variations have been used, among others, to help recover shape and separate diffuse and specular reflectance (e.g., [JA11, ZMKB08]). To the best of our knowledge, chromatic variations have been not exploited to infer reflectance properties such as specular roughness.

3. Shape Agnostic BRDF Estimation

3.1. Input & Assumptions

Our shape agnostic BRDF estimation method takes as input an HDR photograph of the exemplar object and a record of the incident lighting in the form of an HDR light probe. To white balance the acquired HDR photographs, we also place a color checker chart next to the exemplar. Our method relies on matching (binary) color profiles between the captured and rendered images of the candidate reflectance parameters, and as such image noise can produce color artifacts. Therefore, as a preprocessing step, we denoise the input photographs using bilateral filter with a Gaussian kernel, followed by a downsample step that halves the resolution of the photographs.




Figure 2: A red ceramic bowl acquired in an indoor office environment. (a) Material exemplar photograph (HDR). Inset shows the details in the highlights. (b) Corresponding light probe. (c) Diffuse convolution of the probe with an initial estimate of diffuse albedo obtained with the manual selection strategy.

We assume that the exemplar object consists of a homogeneous isotropic dielectric material that can be accurately characterized by the Cook-Torrance microfacet BRDF model [CT82]. However, our method can also be used with any other physically based BRDF model. Furthermore, we assume that the main color in the observed reflection on the dielectric exemplar is due to the diffuse component.

We also assume a single dominant light source in the environment and require it to not be incident at or near a grazing angle to avoid Fresnel gain in the observed reflection of the dominant light. This assumption does not preclude that other light sources can be present in the scene, as long as their contribution is less than that of the dominant light source. In fact, we require surrounding ambient illumination with color variation in the environment illumination. This ensures some chromatic variation in the illumination incident on the exemplar, which we also exploit for robust BRDF estimation in addition to the brightness cue due to the dominant light source. Many outdoor and indoor lighting environments satisfy the above requirements, e.g., lighting coming through a window for an indoor environment, or sunny sky against a colorful outdoor background.

We desire a method suited for non-expert users, and prefer robustness and simplicity over fragile automatic steps. We therefore require a small amount of user direction to bootstrap the algorithm. We ask the user to mark the area of interest on the exemplar photograph for reflectance processing, and optionally also require the user to mark a small window of potentially front-facing pixels on the exemplar for initializing the diffuse albedo estimate (described next).

3.2. Diffuse Initialization

As a first step, we estimate an *initial* estimate of the diffuse albedo. We will refine this estimate in subsequent steps. We provide two alternative strategies: estimation from a manually provided window, and a fully automatic approach. For the manual strategy we rely on a user selected small window of frontal pixels, and simply average the RGB values within this window. For the automatic selection strategy, we sort all exemplar pixels first by intensity (selecting pixels between the lower 15th – 20th percentile), and then based on color saturation (averaging the top 10 percentile) to eliminate potential specular pollution. We found that the manual strategy




Figure 3: Sample candidate parameters of diffuse and specular albedo for varying specular roughness values σ that jointly maximize the match to the maximum intensity of reflection in the input photograph. Inset show the details in the corresponding highlights.

is more robust and requires less iterations to converge, and use this as the default strategy for estimating the diffuse albedo. Figure 2(c) shows the initial diffuse estimation on a sphere for a ceramic red bowl exemplar shown in Figure 2(a).

3.3. Diffuse and Specular Optimization

In the absence of known geometry, there is an inherent ambiguity between the specular albedo and the specular roughness of a material exemplar, as one of these parameters can be altered to compensate for the other along certain metrics, e.g., maximum intensity of observed reflection. We note that the Cook-Torrance BRDF acts as a Gaussian filter, i.e., it averages out the intensity and color of the dominant light source with its surrounding values in the light probe. The higher the specular roughness, the lower the intensity and the color variation in the observed reflection. We make the observation that while we can always alter the specular albedo to compensate for the intensity, this would not affect the color variation seen on the sample. Based on this observation, we jointly tackle the problem of specular albedo and roughness estimation using a two-step approach where we first find the optimal diffuse and specular albedo given a roughness, followed by a roughness selection step.

Albedo Estimation As a precomputation step, we first precompute a diffuse convolution C_d of the light probe assuming a Lambertian BRDF. We also convolve the environment map for specular BRDFs with different roughness values, ranging from 0.005 (i.e., mirror-like) to 0.3 (i.e., near diffuse). We denote each specular convolution as: $C_s(\sigma)$, with σ the roughness value. The resulting reflectance radiance (i.e., rendered sphere) is then:

$$r(\rho_s, \sigma, \rho_d) = \rho_s C_s(\sigma) + \rho_d C_d, \quad (1)$$

with ρ_s and ρ_d the specular and diffuse albedo respectively. For each of the candidate roughness values σ , we solve an optimization problem to find a refined estimate of the diffuse albedo ρ_d and the corresponding specular albedo ρ_s .

However, since we do not know the geometry, we formulate this optimization in terms of *shape agnostic criteria*, namely maximum intensity and color profile matching.

Maximum Intensity Matching If visible, and in absence of noise, the highest intensity on a sphere rendered under the same illumination should be at par with the highest intensity seen in the image I

of the material exemplar:

$$\max(I) = \max(r). \quad (2)$$

In practice, to reduce the impact of residual pixel noise, we average out several (10-20) highest intensity pixels. Given a diffuse estimate ρ_d we can employ this constraint to trivially compute a matching specular albedo ρ_s .

Color Profile Matching The maximum intensity matching only provides a single constraint, and thus cannot solve for both the diffuse and specular albedo. We will therefore only use the maximum intensity constraint to estimate the specular albedo. For the diffuse albedo we will use a color profile matching.

We first define the 3D color histogram where each axis corresponds to the R, G, or B channel. Like any other histograms, bin size is a critical parameter. To ensure a good fillrate, we dynamically find the optimal bin size such that the number of occupied bins is 500. We compute the optimal bin size on the recorded HDR photograph (we denote its histogram h_I), and use it to compute a histogram on the rendered spheres: $h_r(\sigma)$.

To robustly handle the specific factors of matching a profile of a sphere to that of a photograph object of unknown shape (which can have (a) missing normal directions, and (b) different occurrences per normal direction), we cannot rely on standard distribution matching algorithms. Instead, we formulate our color profile matching metric as a combination of two terms. One term matches the volume of the profile:

$$\epsilon_{volume}(\sigma) = |\sum B(h_I) - \sum B(h_r(\sigma))|, \quad (3)$$

where $B(\cdot)$ is the *binary occupancy* operator that returns 1 if a bin is occupied, and 0 when it is empty. The binarization gracefully handles differences in normal occurrence. However, it ignores missing normal directions. The second term measures the importance of mismatched bins (i.e., occupied bins that are empty in the other histogram):

$$\epsilon_{mismatch}(\sigma) = \sum (\delta_{B(h_I) \neq B(h_r(\sigma))} w(h_I, h_r(\sigma))), \quad (4)$$

where $\delta_{x \neq y}$ is the indicator function that is one if one bin is empty and the other is not, and w is an importance weighting function defined as:

$$w(h_I, h_r) = \max(B(h_I)S(h_I, a_I), B(h_r)S(h_r, a_r)), \quad (5)$$

and S is the Sigmoid function:

$$S(x, a) = \frac{1}{1 + e^{-(x-a)}}. \quad (6)$$

Here, we set a_I to be a constant c . We empirically chose $c = 6$ since it sets the range of the Sigmoid to be $(0.0025, 1)$ while does not significantly change the value for the majority of bins. We set a_r to a smaller value than a_I since we trust the entries due to the rendered sphere more than the entries due to the data which suffers from some noise. Hence, we empirically chose a_r to be half of c , and this sets the range of $S(h_r(\sigma), a_r)$ to be $(0.047, 1)$. Intuitively, $\epsilon_{mismatch}$ marks bins in either histogram as less important if the representative color occurs infrequently. This reduces the effects of missing normal directions and noise. However, once a color




Figure 4: RGB space 3D plot of pixels of the exemplar vs. candidate BRDF estimates on a sphere: (a) The red ceramic bowl shown in Fig. 2(a). (b) Sphere rendered with $\sigma = 0.005$ corresponding to the best matching color profile (Fig. 3,a). (c) Sphere rendered with $\sigma = 0.02$ (Fig. 3,b). (d) Sphere rendered with $\sigma = 0.05$ (Fig. 3,c)

is reliably observed, we quickly equalize its weight to avoid over-weighting based on normal occurrence or incident illumination distribution.

We combine these two terms in our final color profile matching loss:

$$\epsilon_{profile}(\sigma) = \epsilon_{mismatch}(\sigma) + \lambda \epsilon_{volume}(\sigma), \quad (7)$$

where λ is a balancing factor between both terms. We found that $\lambda = 0.3$ works well in practice.

We then solve for ρ_s and ρ_d by iterating between solving for each using the maximum intensity matching Equation (2) and color profile matching Equation (7) respectively. We bootstrap this process using the initial diffuse estimate. Figure 3 shows visualization of matched diffuse and specular albedo given a-priori fixed roughness.

Roughness Selection As a final step, we search for the best matching set of overall parameter values (each optimized for a given roughness) by performing a global color profile matching in RGB space of the exemplar data with the best candidate profiles of each specular roughness:

$$\operatorname{argmax}_{\sigma} \epsilon_{profile}(\sigma). \quad (8)$$

As can be seen in Figure 4, the 3D plot corresponding to specular roughness $\sigma = 0.005$ is the best match for the plot of the data from the exemplar photograph, and the final BRDF estimate (along with the final diffuse estimate) can be visually compared against the exemplar in Figure 5. Note that for renderings, we employ the Schlick approximation [Sch94] to model Fresnel gain within the Cook-Torrance BRDF and set the final estimated specular albedo ρ_s as the reflection at normal incidence R_0 for Fresnel gain.




Figure 5: Visual comparison of the red ceramic bowl exemplar (a), and the final optimized diffuse component (b), and the complete BRDF (c).




Figure 6: Appearance estimation of a black leather bag in an indoor environment. (a) Input HDR photograph of the exemplar. (b) Corresponding light probe. (c) A sphere rendered under the same illumination with the estimated BRDF. (d) Rendering with both the estimated BRDF plus synthesized mesostructure details.

4. Synthesizing Mesostructure

The shape agnostic BRDF estimation procedure described in the previous section is agnostic to the underlying shape, and is therefore suitable for estimating the base material BRDF of exemplar objects with both smooth surfaces as well as textured surfaces exhibiting some mesostructure. However, for the latter case, rendering with purely the estimated BRDF is insufficient for reproducing the overall appearance of the material (see Figure 6,c). Hence, for such materials with a texture-like appearance, we additionally synthesize plausible repeating mesostructure details in the form of a surface normal map (to model the shape of the mesostructure) and a specular reflection occlusion map (that models spatially varying specular reflection occlusions inside fine wrinkles and grooves) given the same input HDR photograph of the exemplar object.

For exemplar material samples with a distinctly visible mesostructure, we exploit the visible specular reflection cues (in a local neighborhood) due to the dominant lighting in the surrounding environment to replicate the mesostructure. We synthe-




Figure 7: Target and source tiles for mesostructure synthesis. (a) Target master tile from the black leather bag in Fig. 6. (b) Crop of sphere rendered with estimated BRDF of the black leather exemplar and a borrowed source leather mesostructure from database of Aittala et al. [AWL15]. (c, d) Synthesized surface normals and specular reflection occlusion maps for the target master tile based on best matches found on the source rendered tile.

size not only a “master-tile” of the mesostructure normal-details, but also a specular reflection occlusion map that approximates the reflection occlusions due to fine/deep mesostructures. As we only have a single input photograph, this synthesis process is underconstrained. We bias the synthesis solution by only borrowing mesostructure details from a suitably similar exemplar material in the database of stationary materials previously measured by Aittala et al. [AWL15]. Note that this database provides measurements of spatially varying specular albedo for their reflectance model. Instead, we interpret that data as a spatially varying specular reflection occlusion function that scales the estimated specular albedo of the target material exemplar for modeling the textured appearance along with the shading due to the normal map. While we use the exemplar from this database as a guide, we do *not* require the exemplar to be an exact match for the type of mesostructure seen on our target exemplar material. Instead, we let our synthesis step adapt the source mesostructure details to the structure seen on our target material. This regularization of the synthesis process can be seen as restricting the search for mesostructures to the space spanned by the exemplar’s mesostructures. We posit that similar materials’ mesostructure lie in a similar search space.

Implementation The synthesis step proceeds as follows: we first manually select a relatively flat and frontal facing patch on the exemplar object exhibiting strong specular reflection (due to the dominant light in the lighting environment) as the master tile for which we would like to synthesize the mesostructure details (see Figure 7,a). We then render a sphere with the BRDF estimated for this target exemplar material under the same lighting environment, while borrowing mesostructure details from an appropriate similar looking material in the database (e.g., a leather sample). We select a crop on this rendered sphere around the dominant specular highlight and set that to be the source tile (see Figure 7,b). Then, for every pixel on the target master tile, we search for the best matching pixel (defined by a 15×15 neighborhood) in the source tile and borrow this location’s corresponding surface normal and the spatially varying specular albedo (encoding the specular reflection occlusion) from the known u-v coordinate into the corresponding synthesized mesostructure tiles. For pixels with a partially synthesized mesostructure within their matching neighborhood, we also include the existing surface normals within the neighborhood in the matching function besides the RGB pixel values. This enables

a joint synthesis of both the tiles for surface normals and specular reflection occlusion which is required since they are correlated. As a last step of the synthesis process, we employ the obtained master tiles of mesostructure details to jointly synthesize larger tiles of surface normals and specular reflection occlusion for the final rendering of the exemplar material (Figure 6,d).

5. Results

We now provide additional results for our proposed example-based material appearance estimation method and some analysis of its robustness under various factors. Note that our current implementation is unoptimized and the current timings for the BRDF estimation step is as follows: 1-2 minutes for the initial BRDF estimation procedure using a GPU implementation at 716x716 resolution, and 15-20 minutes for the subsequent optimization using MATLAB on a 6 core (hyper-threaded) desktop computer (Intel Core i7-8700K CPU, 32GB RAM) with an NVIDIA 1080ti GPU.

Qualitative Evaluation We first demonstrate the robustness of our method on a wide variety of materials captured under various indoor and outdoor lighting environments. All examples presented in the paper have been acquired using a Canon 650 DSLR camera with a 50mm lens.

Figure 1 shows an example of a *faux blue leather* purse with a distinct surface mesostructure, acquired at an outdoor park location. Our method is able to obtain a reasonable estimate of the base material BRDF. We also estimate and synthesize appropriate surface mesostructure for the exemplar given another leather mesostructure example. The final rendering with the BRDF estimate and mesostructure faithfully replicates the appearance of the *blue faux leather* material.

Figure 8 presents BRDF estimation results for various exemplar objects with a smooth surface in two different indoor lighting environments. For these examples, the objects' BRDFs range from very sharp specular (*coffee cups* and *plastic cups*) to rough specular (*desktop mouse*) and almost diffuse (*blue fabric*), and from very dark to a bright diffuse albedo. For all of these examples, our method is able to reliably estimate a plausible BRDF fit for the exemplar material that well reproduces the material appearance.

Figure 9 (top-row) presents another example of a *black leather* bag acquired in an indoor environment. This exemplar has slightly larger-scale mesostructure than the *leather bag* example shown in Figure 6 which is appropriately estimated along with the material BRDF. The second and third rows present two different examples of *red* and *light-blue leather* jackets respectively with finer scale mesostructure, that are also estimated appropriately using our synthesis approach. Note that the Aittala et al. database does not include a leather sample with such fine scale mesostructure. Our method is able to successfully adapt the mesostructure of the source exemplar in this case to model the fine-scale mesostructure of these jackets. For such fine scale mesostructure, we employ a smaller matching window of a 3×3 neighborhood for the joint synthesis.

Manual vs. Automatic Diffuse Initialization In section 3.2 we proposed two strategies for obtaining an initial estimate of the diffuse albedo. Figure 10,(b) shows the initially estimated BRDFs using this above automatic diffuse initialization step for the red bowl

Table 1: BRDF fit for the pink glossy ball with our method (shape agnostic) vs. with known shape.

	ρ_d	ρ_s	σ
our method	[0.439, 0.098, 0.136]	0.037	0.07
known shape	[0.542, 0.107, 0.150]	0.025	0.05

Table 2: BRDF fits for the blue faux leather purse under different lighting conditions.

	ρ_d	ρ_s	σ
light probe 1	[0.136, 0.157, 0.224]	0.065	0.2
light probe 2	[0.169, 0.187, 0.275]	0.051	0.2
light probe 3	[0.137, 0.161, 0.204]	0.077	0.2

exemplar shown in Figure 2, the faux blue leather purse shown in Figure 1, and the light-blue leather jacket shown in Figure 9. These initial estimates of the BRDFs are then refined by the subsequent optimization to produce the final BRDF estimates (c) which are consistent with the final BRDF estimates for these exemplars obtained using our default manual initialization of the diffuse albedo (d). Thus, either strategy of automatic or manual initialization can be used in practice with the caveat that we found the automatic initialization to require more computation and iterations for convergence than when employing manual initialization.

Unknown vs. Known Shape Figure 11 shows the BRDF estimated for a *pink glossy plastic* ball in an outdoor environment. Here, the BRDF has been estimated for the ball using our procedure without assuming a spherical shape of the object. As can be seen, the rendered sphere is a close match to the photograph of the ball supporting the quality of the BRDF estimate. For a quantitative comparison, we also fit the BRDF of the pink ball assuming its known spherical shape to obtain a reference. Numerically, the BRDF fit with our shape agnostic procedure is very close to the known shape reference as seen in Table 1.

Different Lighting Figure 12 examines the robustness of our BRDF estimation procedure for the *blue faux leather* purse acquired under three different lighting environments and slightly differing exemplar shape and camera viewpoint. Here, the rendered spheres visualize the obtained BRDF fits compared to the exemplars under the respective illumination conditions. The corresponding numerical fits for the *blue faux leather* purse in these different configurations are given in Table 2. While the BRDF fits vary slightly across the measurements, they are fairly consistent for most of the parameters across the various incident illumination conditions and shape and viewpoint configurations.

Different Mesostructure Exemplar Figure 13 illustrates the impact and importance of selecting an appropriate mesostructure exemplar for the mesostructure synthesis on the *blue faux leather* purse using four different leather exemplars from the Aittala et al. database. The synthesis results for exemplars 1 through 3 are qualitatively similar. This suggests that the synthesis step does not require an exactly matching exemplar in the database and has the ability to adapt the given mesostructure to the master tile. However, the result for exemplar 4 is distinctly degraded compared to the other exemplars, indicating that matching the mesostructure feature scale is important for good qualitative results of the synthesis.



Figure 8: Additional results of example-based material appearance estimation in indoor lighting environments. Top-row: BRDFs estimated in an office environment for two coffee cups with smooth surfaces and a rough specular black plastic mouse. Second-row: BRDFs estimated in a living room environment for a blue plastic cup, a winter jacket, a game controller, and blue apparel fabric.





Figure 9: Additional results of material appearance estimation for exemplar objects with surface mesostructure. Top-row: Black leather bag material estimated in an indoor environment. Second and third rows: a red leather jacket, and a light-blue leather jacket respectively, in an outdoor (backyard) environment.

6. Ablation Study

To gain further insight in the robustness and limitations of our shape-agnostic BRDF estimation, we perform an ablation study using simulated data. We investigate robustness with respect to:

BRDF parameters, exemplar shape, mesostructure, and illumination. Note that we added shot noise to all the simulated exemplar images using the Matlab function *imnoise*.

BRDF Parameters To quantify the robustness with respect to



(a) light probe (b) auto init. (c) final BRDF (d) manual init.

Figure 10: BRDF estimation with automatic diffuse initialization vs. manual initialization for the ceramic red bowl exemplar (top-row), the faux blue leather purse exemplar (center row), and the light-blue leather jacket exemplar (bottom row). (b) Initial BRDF estimate with automatic diffuse initialization. (c) Final BRDF estimate after subsequent optimization. (d) Final BRDF estimate after optimization of manual diffuse initialization.



(a) light probe (b) pink ball (c) est. BRDF

Figure 11: Visual validation of BRDF estimation for a pink glossy plastic ball.

varying specular roughness and various ratios of diffuse and specular albedo, we run our shape-agnostic BRDF estimation method on simulated scenes under a fixed urban lighting environment and a fixed shape (*Utah Teapot*). Table 3 lists the various ground truth parameters and the corresponding estimated parameters. From this we can conclude that accuracy decreases slightly for increasing specular roughness, and for increasing ratio of diffuse to specular albedo. In both cases, the ambiguity in RGB profile matching increases, thereby reducing its accuracy. For visual reference, we also include visualizations for the three least accurate cases from Table 3 (marked in yellow and red) in Figure 14.

Shape Variation We validate the robustness of our method to variations in shape with fixed incident lighting and fixed BRDF parameters (corresponding to *ground truth 2* in Table 3). Figure 15 shows a selection of the various shapes with corresponding recovered BRDFs. As can be seen, the estimated BRDF exhibit very little variations, and are very close to the ground truth, validating the shape-agnostic aspect of our BRDF estimation procedure. We measured the statistics of error in the estimates as follows: ρ_d : RMSE = 0.0061, std. dev. = 0.0021; ρ_s : RMSE = 0.0033, std. dev.



(a) light probe 1 (b) exemplar (c) est. BRDF




(d) light probe 2 (e) exemplar (f) est. BRDF




(g) light probe 3 (h) exemplar (i) est. BRDF

Figure 12: BRDF fits for the blue faux leather purse under different lighting environments.



(a) exemplar 1 (b) exemplar 2



(c) exemplar 3 (d) exemplar 4

Figure 13: Synthesized mesostructure for the blue faux leather purse given four different input leather exemplars. Top-left: Crop of a sphere rendered with estimated BRDF and given leather mesostructure exemplar. Bottom-left: Synthesized surface normals and specular reflection-occlusion of master tile. Right: Sphere rendered with larger synthesized mesostructure tile.

= 0.0029; σ : RMSE = 0.0061, std. dev. = 0.0052. The individual estimated parameters for all the shapes are provided in the supplemental material.

Mesostructure Variation To validate robustness with respect to mesostructure (i.e., high frequency shape variations), we estimate the BRDF with 72 different surface mesostructures from the Ait-tala et al. database mapped to the surface of the *Utah Teapot*. We

Table 3: BRDF estimates for varying parameters of specular roughness and diffuse/specular ratio.

	ρ_d	ρ_s	σ
ground truth 1	[0.3000, 0.0500, 0.0500]	0.060	0.01
estimated 1	[0.3019, 0.0535, 0.0522]	0.061	0.01
ground truth 2	[0.3000, 0.0500, 0.0500]	0.060	0.05
estimated 2	[0.3070, 0.0549, 0.0539]	0.060	0.05
ground truth 3	[0.3000, 0.0500, 0.0500]	0.060	0.10
estimated 3	[0.2886, 0.0407, 0.0386]	0.064	0.10
ground truth 4	[0.5000, 0.1000, 0.1000]	0.040	0.01
estimated 4	[0.5061, 0.1026, 0.1032]	0.041	0.01
ground truth 5	[0.5000, 0.1000, 0.1000]	0.040	0.05
estimated 5	[0.5110, 0.1059, 0.1062]	0.040	0.05
ground truth 6	[0.5000, 0.1000, 0.1000]	0.040	0.10
estimated 6	[0.5260, 0.1185, 0.1180]	0.030	0.07
ground truth 7	[0.7000, 0.2000, 0.2000]	0.020	0.01
estimated 7	[0.7204, 0.2128, 0.2117]	0.019	0.01
ground truth 8	[0.7000, 0.2000, 0.2000]	0.020	0.05
estimated 8	[0.7152, 0.2086, 0.2081]	0.014	0.015
ground truth 9	[0.7000, 0.2000, 0.2000]	0.020	0.10
estimated 9	[0.7265, 0.2137, 0.2148]	0.016	0.07




Figure 14: Visualizations for the three least accurate cases recovered in the BRDF parameter ablation study in Table 3.

use the same BRDF as in the *Shape ablation study*. We found the BRDF estimates to be very close to the ground truth values for all but 8 of the tested mesostructures which are shown in Figure 16. These mesostructures exhibit strong directional structures and/or anisotropy which violates our assumptions of isotropic surface appearance. These 8 examples have been identified based on the numerical difference in the estimated parameters to the ground truth being above a set threshold (provided in the supplemental material). However, we note that only two of these examples (1 and 6) exhibit a perceptually noticeable difference in the BRDF fit. Thus,




Figure 15: A selection of shapes explored for the Shape ablation study. The resulting recovered BRDFs are a close match, indicating that our method is robust to shape variations.

we infer that surface mesostructure has very little (if any) effect on the BRDF estimation.

Illumination Variation To validate the robustness to variations in illumination, we estimate the BRDF (same as in the *Shape ablation study*) on the *Utah Teapot* shape under 8 different lighting environments (3 indoors and 5 outdoors) shown Figure 17. As can be seen, the BRDF estimation is consistent across the different lighting environments, and very close to the ground truth values for all cases except for the Uffizi gallery example (bottom). The specular roughness is overestimated in the case of Uffizi due to lack of color variation in the lighting environment (overcast sky, gray walls) which increases ambiguity for the RGB color profile matching. In all other cases, there is sufficient natural color variation in the lighting environment for the method to work robustly. We measured the statistics of error in the estimates as follows: ρ_d : RMSE = 0.0072, std. dev. = 0.0043; ρ_s : RMSE = 0.0043, std. dev. = 0.0045; σ : RMSE = 0.0094, std. dev. = 0.0099. The individual estimated BRDF parameters under each of these lighting environments is provided in the supplemental material.

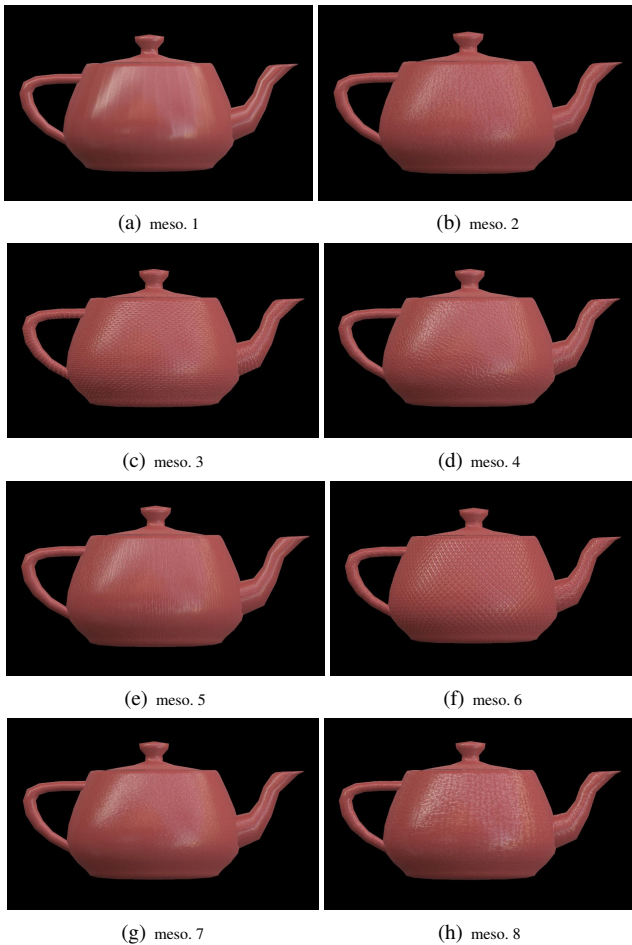


Figure 16: Visualizations of the eight least accurate cases of recovered BRDFs in the Mesostructure variation ablation study conducted using 72 measured mesostructures in the Aittala database.

7. Limitations

Our shape-agnostic BRDF estimation procedure relies on binary RGB color profile matching for BRDF estimation. When the incident lighting exhibits insufficient color variation, e.g., on a very overcast day with gray sky, or colorless surroundings (white/gray walls, streets), our method is unable to distinguish the different BRDFs. Empirically, we found that lighting environments with moderate dynamic range variation in illumination incident from different directions works best. For example, such condition occur naturally in outdoor environments where the sky and the ground adds sufficient background illumination to the dominant sunlight. In indoor environments, the method works better in situations where the room is well lit with bounce light from all sides besides the dominant light (e.g, from a window), and fails when the only incident illumination is due to a dominant source and the rest of the incident directions are relatively dark. Furthermore, the accuracy of the apparent BRDF is also affected by the frequency of the incident illumination and the scale of observation of the exemplar; this limitation is inherent and shared by most methods that estimate surface



Figure 17: A selection of illumination environments explored for the Illumination ablation study. The resulting recovered BRDFs are a close match to ground truth for all cases except the Uffizi gallery environment (bottom).

reflectance from natural lighting. Finally, the mesostructure synthesis relies on a suitable mesostructure exemplar and can fail when the source mesostructure is very dissimilar to the target mesostructure, or in the absence of sufficient specular cues in the reflection for guiding the synthesis. The synthesis step currently employs an unoptimized CPU implementation and could be significantly sped up with a GPU implementation.

8. Conclusion

We presented a novel example-based “rapid-appearance-modeling” technique for homogeneous isotropic dielectric objects suitable for on-site capture by non-expert users. A key distinguishing feature of our method is that we do not require prior knowledge of the shape

of the exemplar object; we only require knowledge of the incident lighting which is simpler to acquire. We introduce two shape agnostic constraints (i.e., maximum intensity and color profile matching) that guide the optimization process of the reflectance parameters. In addition, we also model the appearance of materials exhibiting a regular or stationary mesostructure by jointly synthesizing appropriate surface normals and spatially varying specular reflection occlusion for modeling the target material's mesostructure.

Currently, our process involves a few easy manual steps which are well suited for digital artists. Automating these steps is an interesting avenue for future work, especially in biasing the mesostructure synthesis by material type. Extending the approach for modeling spatial variation in the appearance of objects with arbitrary shapes is also interesting future research direction.

9. Acknowledgements

This work was supported by an EPSRC Early Career Fellowship (EP/N006259/1) and a GPU donation from Nvidia to Abhijeet Ghosh. Pieter Peers was partially supported by NSF grant IIS-1350323 and gifts from Google, Activision, and Nvidia.

References

- [AAL16] AITTALA M., AILA T., LEHTINEN J.: Reflectance modeling by neural texture synthesis. *ACM Trans. Graph.* 35, 4 (July 2016), 65:1–65:13. 1, 2
- [AWL15] AITTALA M., WEYRICH T., LEHTINEN J.: Two-shot svbrdf capture for stationary materials. *ACM Transactions on Graphics* 34, 4 (2015), 110. 1, 2, 5
- [BJK07] BASRI R., JACOBS D. W., KEMELMACHER I.: Photometric stereo with general, unknown lighting. *IJCV* 72, 3 (2007), 239–257. 2
- [BM15] BARRON J. T., MALIK J.: Shape, illumination, and reflectance from shading. *IEEE PAMI* (2015). 2
- [CGS06] CHEN T., GOESELE M., SEIDEL H. P.: Mesostructure from specularities. In *CVPR* (2006), pp. 1825–1832. 2
- [CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM TOG* 1, 1 (1982), 7–24. 2, 3
- [DAD*18] DESCHAIANTRE V., AITTALA M., DURAND F., DRETTAKIS G., BOUSSEAU A.: Single-image svbrdf capture with a rendering-aware deep network. *ACM Trans. Graph.* 37, 128 (Aug. 2018), 15. 2
- [DCP*14] DONG Y., CHEN G., PEERS P., ZHANG J., TONG X.: Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting. *ACM Trans. Graph.* 33, 6 (2014), 193:1–193:12. 2
- [DRS08] DORSEY J., RUSHMEIER H., SILLION F.: *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., 2008. 2
- [GRR*17] GEORGIOULIS S., REMATAS K., RITSCHER T., GAVVES E., FRITZ M., GOOL L. V., TUYTELAARS T.: Reflectance and natural illumination from single-material specular objects using deep learning. *PAMI* (2017). 2
- [HS03] HERTZMANN A., SEITZ S. M.: Shape and materials by example: A photometric stereo approach. In *CVPR* (2003), pp. 533–540. 2
- [HSL*17] HUI Z., SUNKAVALLI K., LEE J., HADAP S., SANKARANARAYANAN A.: Reflectance capture using univariate sampling of brdfs. In *ICCV* (2017). 1, 2
- [JA11] JOHNSON M. K., ADELSON E. H.: Shape estimation in natural illumination. In *CVPR* (June 2011), pp. 2553–2560. 2
- [LDPT17] LI X., DONG Y., PEERS P., TONG X.: Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Trans. Graph.* 36, 4 (July 2017), 45:1–45:11. 2
- [LN16] LOMBARDI S., NISHINO K.: Reflectance and illumination recovery in the wild. *IEEE PAMI* 38, 1 (2016), 129–141. 2
- [LSC18] LI Z., SUNKAVALLI K., CHANDRAKER M. K.: Materials for masses: Svbrdf acquisition with a single mobile phone image. 2
- [LXR*18] LI Z., XU Z., RAMAMOORTHY R., SUNKAVALLI K., CHANDRAKER M.: Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM Trans. Graph.* 37, 6 (Dec. 2018). 2
- [MMZ*18] MEKA A., MAXIMOV M., ZOLLHOEFER M., CHATTERJEE A., SEIDEL H.-P., RICHARDT C., THEOBALT C.: Lime: Live intrinsic material estimation. In *CVPR* (June 2018). 2
- [ON16] OXHOLM G., NISHINO K.: Shape and reflectance estimation in the wild. *PAMI* 38, 2 (Feb. 2016), 376–389. 2
- [PCDS12] PALMA G., CALLIERI M., DELLEPIANE M., SCOPIGNO R.: A statistical method for svbrdf approximation from video sequences in general lighting conditions. *Comput. Graph. Forum* 31, 4 (2012), 1491–1500. 2
- [RPG16] RIVIERE J., PEERS P., GHOSH A.: Mobile surface reflectometry. *Computer Graphics Forum* 35, 1 (2016), 191–202. 1, 2
- [RVZ08] ROMEIRO F., VASILYEV Y., ZICKLER T.: Passive reflectometry. In *ECCV* (2008), pp. 859–872. 2
- [RZ10] ROMEIRO F., ZICKLER T.: Blind reflectometry. In *ECCV* (2010), pp. 45–58. 2
- [Sch94] SCHLICK C.: An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246. 4
- [THS04] TREUILLE A., HERTZMANN A., SEITZ S. M.: Example-based stereo with general BRDFs. In *ECCV* (2004), pp. 457–469. 2
- [WK15] WEINMANN M., KLEIN R.: Advances in geometry and reflectance acquisition. In *ACM SIGGRAPH Asia, Course Notes* (2015). 2
- [Woo80] WOODHAM R. J.: Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19, 1 (1980), 3050–3068. 2
- [WSM11] WANG C.-P., SNAVELY N., MARSCHNER S.: Estimating dual-scale properties of glossy surfaces from step-edge lighting. *ACM Trans. Graph.* 30, 6 (2011), 172:1–172:12. 2
- [WWZ16] WU H., WANG Z., ZHOU K.: Simultaneous localization and appearance estimation with a consumer rgb-d camera. *IEEE Transactions on Visualization and Computer Graphics* 22, 8 (Aug 2016), 2012–2023. doi:10.1109/TVCG.2015.2498617. 2
- [XDPT16] XIA R., DONG Y., PEERS P., TONG X.: Recovering shape and spatially-varying surface reflectance under unknown illumination. *ACM Transactions on Graphics* 35, 6 (December 2016). doi:https://doi.org/10.1145/2980179.2980248. 2
- [XNY*16] XU Z., NIELSEN J. B., YU J., JENSEN H. W., RAMAMOORTHY R.: Minimal brdf sampling for two-shot near-field reflectance acquisition. *ACM Trans. Graph.* 35, 6 (Nov. 2016), 188:1–188:12. 2
- [YLD*18] YE W., LI X., DONG Y., PEERS P., TONG X.: Single photograph surface appearance modeling with self-augmented CNNs and inexact supervision. *Computer Graphics Forum* 37, 7 (Oct 2018). doi:https://doi.org/10.1111/cgf.13560. 2
- [ZCD*16] ZHOU Z., CHEN G., DONG Y., WIPF D., YU Y., SNYDER J., TONG X.: Sparse-as-possible svbrdf acquisition. *ACM Trans. Graph.* 35 (November 2016). 2
- [ZMKB08] ZICKLER T., MALLICK S. P., KRIEGMAN D. J., BELHUMEUR P. N.: Color subspaces as photometric invariants. *IJCV* 79, 1 (Aug 2008), 13–30. 2