

BSSRDF Estimation from Single Images

Adolfo Munoz¹ Jose I. Echevarria¹ Francisco J. Seron¹ Jorge Lopez-Moreno¹ Mashhuda Glencross² Diego Gutierrez¹

¹Universidad de Zaragoza ²Loughborough University

Abstract

We present a novel method to estimate an approximation of the reflectance characteristics of optically thick, homogeneous translucent materials using only a single photograph as input. First, we approximate the diffusion profile as a linear combination of piecewise constant functions, an approach that enables a linear system minimization and maximizes robustness in the presence of suboptimal input data inferred from the image. We then fit to a smoother monotonically decreasing model, ensuring continuity on its first derivative. We show the feasibility of our approach and validate it in controlled environments, comparing well against physical measurements from previous works. Next, we explore the performance of our method in uncontrolled scenarios, where neither lighting nor geometry are known. We show that these can be roughly approximated from the corresponding image by making two simple assumptions: that the object is lit by a distant light source and that it is globally convex, allowing us to capture the visual appearance of the photographed material. Compared with previous works, our technique offers an attractive balance between visual accuracy and ease of use, allowing its use in a wide range of scenarios including off-the-shelf, single images, thus extending the current repertoire of real-world data acquisition techniques.

Categories and Subject Descriptors (according to ACM CCS): I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Reflectance I.3.3 [Computer Graphics]: Picture/Image Generation—Display Algorithms I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Shading

1. Introduction

Rendering algorithms have evolved considerably over the past decades, which in turn has motivated new acquisition methods of reflectance data from real-world objects. While this is still an active area of research [WLL*08,GJJD09], the ability to estimate the reflectance characteristics of materials from a single image remains a considerable challenge. Given sparse photographic input, it is impossible to infer the exact geometry and lighting captured in a photograph, which are necessary for an accurate capture. Thus, additional hardware and multiple images are usually employed to obtain that information.

In this work, we present a method to obtain an *approximation* of the Bidirectional Subsurface Scattering Reflectance Distribution Function (BSSRDF) of translucent, homogeneous objects from a single image, based on the diffusion approximation [JMLH01]. Under unknown lighting conditions and assuming no previous knowledge of the scene, this is a very ill-posed problem, which makes it impossible to recover

the exact BSSRDF. Our goal is to devise a simple image-based capture algorithm which yields a physically plausible function that captures the appearance of the material and can be used for rendering. Figures 1, 9 and 10 show some of our results.

Our algorithm is robust and captures accurate BSSRDFs under controlled conditions, compared against physically measured data. We test its robustness in uncontrolled environments, where neither the lighting nor the geometry in the image are known. To approximate these, we extend existing techniques to the more difficult case of translucent materials. We show that even in the presence of such suboptimal input, our estimated BSSRDF succeeds at capturing the visual appearance of such materials.

2. Previous Work

A wide range of methods for measuring reflectance properties from real-world samples exists. These typically use



Figure 1: Starting with a single image, and without any other prior information, we capture an approximation of the subsurface scattering properties of objects with varying degrees of translucency. Then, we use the estimated BSSRDFs to render objects made of similar materials. From left to right: grape, orange soap and wax. The source photos are shown in the insets.

specialized equipment such as a gonireflectometer and/or photographic input obtained over a range of known viewing and lighting directions, e.g. [LKG*03, ST06]. Single image approaches that require prior knowledge about the shape of the object have also been developed [BG01]. These methods usually aim at capturing a representation of the BRDF of opaque objects; we refer the reader to the excellent existing literature for a more comprehensive description [DRS07, WLL*08].

Capturing and modeling the BSSRDF of *translucent* materials is a harder problem that generally requires the use of special measuring setups and long capture sessions (see for instance [JMLH01, GLL*04, WMP*06]). Camera-projector systems have also been used to measure reflectance of small material samples [PvBM*06, TGL*06]. More recent approaches aim to capture BSSRDF models using more practical camera equipment. Donner and colleagues [DWD*08] use multi-spectral images to measure skin reflectance, requiring samples to be taken in front of their capture setup. Another approach exploits cross-polarization photography and uses 20 photographs from a single viewpoint to acquire a layered reflectance model of skin [GHP*08]. The final example in this kind of approaches requires sampling a cube of the material to be captured, constraining the position of the camera and light source [WZT*08]. Other alternative approaches aim to separate the subsurface scattering component of objects in an image, either by adding a set of diffuse priors [WT06] or using high-frequency patterns of illumination in a set of images [NKGR06]. No specific reflectance model parameters are estimated, and thus using the results in a different context remains an open problem. The recently published SubEdit system [STPP09] includes the possibility of hallucinating a BSSRDF from two inputs: a single photograph under fixed lighting, plus previously acquired data from one or more different BSSRDFs. The user assigns scattering profiles from the measured data set to representative points in the image, and the effect is propagated across the surface. Our approach does not require the user to mark corresponding scattering functions and does not require the use of previously measured data. The Lit Sphere user-guided ap-

pearance transfer approach [SMGG01] transfers shading information from an image of a lit sphere to a complex object. In contrast to our work, this approach requires user interaction and would not allow relighting of the original material. It is also unclear how such approximation could be extended for translucent materials.

Recently, there have been two works that focus on estimating translucency properties from single images [MSY09, MMTG09]. Both propose methods that approximate scattering properties of objects under controlled settings, based on the dipole approximation. In contrast to our approach, they require the 3D location of the camera, the lighting configuration of the scene and the geometry of the target object to be known a-priori. Additionally, the method by Mukaigawa et al. [MSY09] require the use of manually-rotated polarizing filters and light-absorbing black sheets during the capture. As acknowledged in their paper, their approach is quite unstable despite this dedicated hardware; this limits the applicability of the method, as their reduced set of results suggests.

3. BSSRDF Estimation

Our BSSRDF estimation is based on the diffusion approximation [JMLH01] and is performed in two steps. First, the diffusion profile is expressed as a linear combination of piecewise constant basis functions, resulting in a linear system that can be efficiently solved applying the Quasi-Minimal Residual method [BBC*94]. This increases the robustness of the method in the presence of suboptimal input derived from our ill-posed, uncontrolled scenarios. The second step performs a smoothing over the profile, eliminating discontinuities on the first derivative and ensuring physical plausibility. A reasonable option for the algorithm design would be a single-step non-linear optimization. However, our preliminary tests [MMTG09] show that due to the ill-posed and underconstrained nature of the problem, this usually reaches local minima, yielding no plausible results.

In this section we introduce our approach for controlled environments, where both the geometry and the main light

	Piecewise constant			Piecewise linear [MSY09]			Zero-mean gaussian			Hermite polynomials			Legendre polynomials		
Number of functions	10	20	30	10	20	30	10	20	30	10	20	30	10	20	30
Estimation time	24 s	31 s	37 s	32 s	45 s	59 s	10 m	20 m	31 m	86 s	-	-	91 s	198 s	14 m
Condition number	$2.9 \cdot 10^3$	$2.2 \cdot 10^4$	$5.9 \cdot 10^4$	$1.6 \cdot 10^4$	$1.9 \cdot 10^6$	$6.3 \cdot 10^6$	$1.7 \cdot 10^7$	$7.4 \cdot 10^7$	$2.3 \cdot 10^8$	$1.9 \cdot 10^{12}$	-	-	$4.6 \cdot 10^7$	$2.1 \cdot 10^9$	$1.9 \cdot 10^{10}$
Error	1.13	0.69	0.86	2.70	1.86	5.24	134.23	356.84	1708.74	3.47	-	-	5.23	5.41	4.85

Table 1: Results from our basis functions tests for the skull made of whole milk material [JMLH01] from Figure 3. For an increasing number of basis functions, the table shows estimation time, condition number of the matrices and error of the resulting diffusion profile (defined as $\int_0^1 [R_d(r) - \sum_{h=1}^m \hat{w}_h B_h(r)]^2 dr$, where R_d is the original diffusion profile). For more than 20 Hermite polynomials the system does not converge. For piecewise linear representation, the first row refers to the number of points of the piecewise linear representation.

Equation 3 for each pixel $i \in \mathbb{I}$. We first rewrite Equation 3 as:

$$L_i = \sum_{j \in \mathbb{O}} (K_j R_d(r) + K_{b,j} R_b(r_b)) \quad (4)$$

where $K_j = \pi^{-1} F_t(\eta, c) E_j \Delta A$ (with a similar definition for $K_{b,j}$). Next, we estimate R_d by a linear combination of m basis functions:

$$R_d(r) \approx \sum_{h=1}^m \hat{w}_h B_h(r) \quad (5)$$

where $B_h(r)$ represents the basis functions (discussed at the end of this section) and \hat{w}_h are the weights assigned to each basis function. Equation 4 now yields:

$$L_i = \sum_{j \in \mathbb{O}} \left(K_j \sum_{h=1}^m \hat{w}_h B_h(r) + K_{b,j} \sum_{h=1}^m \hat{w}_h B_h(r_b) \right) \quad (6)$$

This equation applies to every pixel $i \in \mathbb{I}$, so the complexity of this algorithm is $O(p^2)$ (where p is the number of pixels of the image). However, we have found that down-scaling \mathbb{I} to a resolution of around 200x200 (preserving the aspect ratio of the input image) yields valid approximations for R_d while greatly reducing computation times. Applying the equation to each pixel of the scaled \mathbb{I} we get a linear system defined by the matrix product $\mathbf{A} \cdot \mathbf{X} = \mathbf{B}$, for n pixels and m basis functions, with:

$$a_{ih} = \sum_{j \in \mathbb{O}} (K_j B_h(r) + K_{b,j} B_h(r_b)) \quad (7)$$

$$\mathbf{X}_{m \times 1}^T = (\hat{w}_1 \quad \hat{w}_2 \quad \dots \quad \hat{w}_m) \quad (8)$$

$$\mathbf{B}_{n \times 1}^T = (L_1 \quad L_2 \quad \dots \quad L_n) \quad (9)$$

Resolution method: To solve the equivalent system $(\mathbf{A}^T \mathbf{A}) \mathbf{X} = (\mathbf{A}^T \mathbf{B})$ we note that some columns in \mathbf{A} may contain values close to zero. This leads to a highly ill-conditioned matrix, while the related basis functions have negligible influence in the final solution. We thus set the associated weights \hat{w}_h to 0 and remove the corresponding columns from \mathbf{A} . Although this approximation reduces the condition number, the system is still ill-conditioned; we improve it further by using a Jacobi pre-conditioner for $(\mathbf{A}^T \mathbf{A})$, and solve the system using the Quasi-Minimal Residual (QMR) method [BBC*94].

Basis functions: In order to choose an appropriate set of basis functions, we rendered translucent objects using measured materials [JMLH01]: in their work, the authors obtain scattering parameters by illuminating the surface of a translucent sample with focused white light and photograph it using a 3-CCD video camera. We then used the resulting renderings along with known geometry and lighting as input to approximate their diffusion profiles testing different options: uniformly distributed piecewise constant functions, zero-mean gaussians (inspired by the work of d'Eon et al [dLE07]), Hermite and Legendre polynomials. Another option that has been previously used to represent diffusion profiles are piecewise linear polynomials [XGL*07, MSY09]. To be able include them in our tests we use the more recent formulation by Mukaigawa et al [MSY09].

Zero-mean gaussian functions, Hermite and Legendre polynomials show high condition numbers, thus leading to unstable linear systems (see Table 1). Hermite polynomials do not even converge for 20 basis functions or more, while gaussian functions show very high errors. On the other hand, the condition number of piecewise linear functions [MSY09] is two orders of magnitude higher, and the error between two and six times larger than piecewise constant functions, which show the best overall behavior while being the fastest to compute. We thus choose to represent diffusion profiles with these basis functions in the first step of our algorithm. A good compromise between detail in the estimation and system stability is reached by using between 20 and 30 basis functions.

This difference between the stability of piecewise constant functions and the other presented options becomes very relevant in the case of inaccurate inputs, which is always the case when generalizing to uncontrolled single images (see Section 4). We found that, in those cases, more unstable bases such as Legendre polynomials or piecewise linear functions lead to higher condition numbers and the QMR method does not often converge to a solution.

Smoothing: In our second step, we fit this piecewise constant profile to a continuous, differentiable, monotonically decreasing function. This helps to eliminate noise and avoid discontinuities in the renderings, while keeping the function physically plausible. Our algorithm does not impose a particular model for this function, although the logical option would be to fit both scattering and absorption of the dipole

	Reduced albedo									Reduced extinction (mm^{-1})								
	[JMLH01]			Estimated			Error			[JMLH01]			Estimated			Error		
	R	G	B	R	G	B	R	G	B	R	G	B	R	G	B	R	G	B
Apple	0.9987	0.9986	0.9772	0.9969	0.9985	0.9686	0.18%	0.01%	0.88%	2.2930	2.3934	2.0160	2.2428	2.3216	2.0202	2.19%	3.00%	0.21%
Cream	1.0000	0.9995	0.9949	1.0000	1.0000	0.9967	0.00%	0.05%	0.18%	7.3802	5.4728	3.1663	7.4580	5.9233	3.4267	1.05%	8.23%	8.22%
Marble	0.9990	0.9984	0.9976	1.0000	1.0000	1.0000	0.10%	0.16%	0.24%	2.1921	2.6241	3.0071	2.3543	2.7351	3.0359	7.40%	4.23%	0.96%
Potato	0.9965	0.9873	0.8209	1.0000	0.9999	0.9145	0.35%	1.27%	11.40%	0.6824	0.7090	0.6700	0.6690	0.6806	0.5651	1.97%	4.00%	15.65%
Skim milk	0.9980	0.9980	0.9926	0.9898	1.0000	0.9981	0.82%	0.20%	0.56%	0.7014	1.2225	1.9142	0.6875	1.2602	1.8943	1.99%	3.08%	1.04%
Whole milk	0.9996	0.9993	0.9963	1.0000	1.0000	0.9818	0.04%	0.07%	1.46%	2.5511	3.2124	3.7840	2.4968	3.1725	3.7553	2.13%	1.24%	0.76%

Table 2: Comparison between the measured properties of several materials [JMLH01] and the estimated properties resulting from our method, fitted to the dipole model.

model [JMLH01]. However, working with a single image, it is not possible to deduce the physical size of the object nor the power of the light source, both necessary to obtain the corresponding dipole diffusion profile.

Thus, we propose a piecewise cubic polynomial $\hat{R}_d(r)$ instead, using Hermite interpolation. This model is generic and not associated to any physically-based BSSRDF model, which makes the method more flexible. The set of points and derivatives of this function is obtained by using a Simulated Annealing algorithm to minimize the following energy function:

$$E = w_d \int_0^1 \left(\hat{R}_d - \sum_{h=1}^m w_h B_h \right)^2 dr + w_p \int_0^1 \left(\hat{R}_d' \right)^2 \delta_{R_d'}(\mathbb{R}^+) dr + w_s \int_0^1 \left(\hat{R}_d'' \right)^2 dr \quad (10)$$

where δ represents the Dirac measure function and w_d , w_p and w_s represent the weights of each term (which we experimentally set to 1, 10 and 10^{-4} , respectively). The first term is related to the difference between the smoothed function and the linear combination; the second term preserves the physical plausibility of the profile by penalizing positive derivatives, and the third term preserves the smoothness of the function. The dependencies on r have been omitted for the sake of clarity.

Validation: In order to validate our BSSRDF estimation algorithm independently of the accuracy of the input data, we first test it under known geometry and lighting (which allows us to use the dipole model): we again rendered objects with different measured material parameters [JMLH01] and then used the resulting images as input to our algorithm. To derive reduced albedo and extinction coefficients and thus provide an accurate numerical comparison, the estimated piecewise constant diffusion profiles were fitted in this case to the dipole model. Note that, as stated before, this fitting to the dipole is not possible for uncontrolled environments, and is introduced here for validation purposes only. For the rest of the paper, we use the piecewise cubic polynomial previously introduced.

Table 2 compares our results with the original physically measured data [JMLH01]; it can be seen how our method yields very small residual error for most materials. As a result, both the profiles and the overall look of the images rendered with them are very similar to the ground truth (see Figure 3). The differences are due to the coarse modeling of the R_d function by a limited number of basis functions, given the intrinsic trade-off between this number and the conditioning of the linear system.

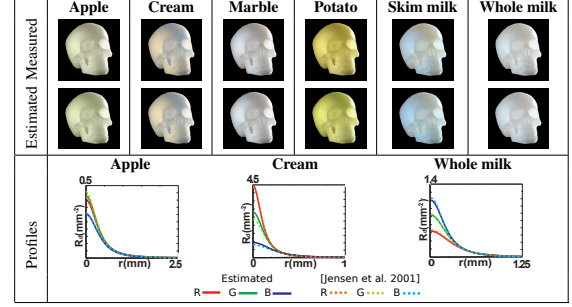


Figure 3: Top two rows: Comparison between renderings using physically measured materials [JMLH01] and our estimated diffusion profiles. Bottom row: Comparison of diffusion profiles. Please refer to the supplementary material for the whole set of profiles.

4. Estimation from Uncontrolled Single Images

We have demonstrated the suitability of our method in controlled environments. In this section we extend our approach to a much more challenging scenario: approximating diffusion profiles from uncontrolled single images. This is a heavily ill-posed problem, given that neither the light direction nor the geometry are known in this case. Therefore, instead of trying to recover an exact physically-based BSSRDF (which is obviously impossible), we aim to estimate a plausible representation that yields results similar to the material depicted in the input image.

We leverage the findings by Fleming and colleagues [FB05], who conclude that humans do not understand translucency through accurate inverse optics, but instead perceive the overall look of translucent materials based on simple image heuristics. This suggests that a suitable *approximation* of both the shape of the object and incident light direction may suffice for our purposes. We extend the usability of existing techniques, originally devised for opaque objects, and show that they can still yield plausible results when complying with our initial assumptions of global convexity and distant light sources.

Estimating shape: Estimating shape from a single image of an opaque object is an under-constrained problem by itself. Previous works, however, have shown how rough approximations can work well in the context of material editing [KRFB06] or the simulation of caustics [GSLM*08]. We note that this estimation is even harder if the object is translucent, given the softening effects of subsurface scattering; we

aim to find a similar approximation that works well for our purposes.

We base our estimation on three sources of information: pixels in the contour (which we assume to lie on the image plane at $Z=0$), shading information across its surface and the assumption of global convexity [LB00]. Inspired by previous approaches [KRFB06, Joh02], we reconstruct the depth map Z of an object as the weighted sum of a base layer (which encodes global convexity) and a detail layer (which encodes high frequency), both obtained by means of the bilateral filter. We use values of $\sigma_{spatial} \in [0.08..0.1]$ and $\sigma_{intensity} \in [0.3..0.5]$ for the bilateral filter, while the weights for adding the base and detail layers are usually 0.8 and 0.2 respectively (thus favoring global convexity over details). We rely on additional non-linear spline functions to reshape the base layer and boost its apparent "inflation" [KRFB06]. Given the inherent bass-relief ambiguity, we reverse the resulting signal if necessary to comply with our global convexity assumption, which yields our final depth map Z . A normal map N is subsequently computed from Z . Additionally, a back depth map Z_b plus the corresponding back normal map N_b are generated. We make the simplifying assumption that the back of the object can be approximated by mirroring Z . While this is a strong simplification to circumvent the fact that we do not have information about the back portion of the object in the image, this straightforward operation suffices to produce good results when the object is not strongly illuminated from its back side. In fact, note that the heart-shaped soaps from Figure 10 and the mouse-shaped soap from Figure 1 are not symmetrical (their back face is plain) but still yield plausible profiles.

It could be argued that a simpler depth-recovery technique could be used instead, but in our experiments (which can be found in the supplementary material[‡]) this approach showed a good compromise between quality of the results and ease of use. We nevertheless restrict our estimations to simple geometries in order to minimize the impact of this error on the BSSRDF estimation, leaving the field of depth estimation from complex translucent geometries still open for further research. In the future, more accurate techniques could be trivially included at this stage.

Estimating light direction: Several existing methods can estimate light source directions from a single image, but usually at the expense of assuming some previous knowledge or including a calibration object in the scene [ZY01, WS02]. In contrast, our goal is to obtain the dominant light direction starting with a single, off-the-shelf image, and thus we cannot impose such restrictions to our inputs.

We apply the method recently proposed by Lopez-Moreno et al. [LMHRG10], which performs a two-step analysis of the luminance channel of an object: first, the pixels of the

contour \mathbb{O}' are clustered by a k-means algorithm to identify the number of light sources in the scene, as well as their azimuth θ_i direction (in image-space) and relative intensities. Second, zenith angles ϕ_i are approximated for each light direction by analyzing gradients in the interior of the object. The pair (θ_i, ϕ_i) defines the recovered 3D direction for each light.

Note that the original light detection algorithm was designed for opaque objects. In order to assess how well it extends to translucent objects, we tested it in controlled scenes with incident lights at specific directions over different objects with varying degrees of translucency. In our tests with different degrees of translucency, the error of the algorithm was always less than 20° , which has been found to be below perceptual threshold [LMSSG10]. The complete test with the different geometries, levels of translucency and light positions, plus another test of the behavior of the BSSRDF estimation algorithm when the input light directions are not accurate, can be found in the supplementary material.

Size of the object: Automatic estimation of the actual size of an object from a single photograph is not possible. Given that the diffusion profile R_d is a function of distance, we use a normalized unit distance equal to the width of the object in the image, and distribute all the piecewise constant basis functions in the range $[0, 1]$. In order to change the relative apparent size of the new rendered objects, it is possible to scale the diffusion profile as follows [STPP09]:

$$R'_d(r) = \frac{1}{s^2} R_d\left(\frac{r}{s}\right) \quad (11)$$

where s is the scaling factor and R'_d is the scaled diffusion profile. Figure 4 shows the effect of this scaling.

5. Results and Discussion

Figure 5 shows the complete validation of the whole pipeline. We first rendered a heart-shaped object with three different measured materials (potato, marble and apple). We then used the rendered images as the only input to our algorithm (no geometry nor lighting are known) and approximated the BSSRDF from them. Finally, we re-rendered the same object with the resulting function. As it can be seen, the estimated materials achieve a very good visual match when compared to the original renderings.

We solve Equation 6 in 30-40 seconds on a Dual Opteron @2.2 GHz with 4 GB of RAM, using between 20 and 30 basis functions for our representation. The smoothing step takes around 20 additional seconds. The recovered BSSRDF for the different materials can be directly used for rendering with no restrictions: for different geometries and under different illumination conditions. Figures 1 and 10 show several results for a wide range of translucent materials, including wax, soap, milk, ketchup, orange juice, detergent, grape

[‡] http://giga.cps.unizar.es/~amunoz/projects/EG2011_bssrdf




Figure 4: Relative sizes for the same material. Left: blue soap from Figure 10. Right: grape from Figure 1

and human skin[§]. Our method works well even for extremely complex materials like skin, although it obviously cannot reproduce the subtleties of light transport in its multi-layered structure. Note that the renderings include additional specular highlights (Phong model) not captured with our method. The lighting in those figures has been set up to match the source image for easy direct comparison: more results under different lighting conditions and geometries can be seen in Figure 9 or the supplementary video, and with different relative sizes for the same material in Figure 4.




Figure 5: Validation of the whole algorithm. Top row: render of measured materials [JMLH01]. Bottom row: our resulting estimations without any prior information. From left to right: potato, marble and apple.

As the results show, our method is fairly robust to inaccurate inputs, although it presents some limitations. In the case of uncontrolled images, large errors in depth or light estimations[¶] may of course lead to larger errors in the results. We are, therefore, bound by the current state of the art in depth and light approximation algorithms from single images, which in practice means that the algorithm works better with images showing simple, convex shapes lit from

[§] Please refer to the supplementary material for higher-resolution versions, along with the mattes, depth, recovered diffusion profiles and light directions for each material.

[¶] Examples can be found in the supplementary material




Figure 6: Comparison between our estimated potato material from Figure 5 and the source potato material from [JMLH01]. Left column: Applying both materials to a new geometry. Middle column: Applying both materials to new geometry under new illumination conditions. Right column: Larger size of the geometry

one direction. Furthermore, our approximation of the geometry of the back side prevents us from estimating the material from objects that present a strong illumination from its back.

Our algorithm works only with the information that is present in the source image. It is therefore expected to be less accurate with sub-optimal input data when estimating parts of the diffusion profile that are not represented in the source image and thus sub-optimally represented in the captured profile. Figure 6 (top row) shows our captured potato material from Figure 5 rendered over different geometry and light directions; the bottom row depicts the equivalent results using the physically measured material [JMLH01] for comparison purposes. Our algorithm, handles this lack of information pretty well when geometry or lighting change substantially from the original image. However, when the size of the geometry changes, the final rendering may deviate from the ground truth reference, as the render is accessing parts of the diffusion profile that were not represented in the source image. Nevertheless, the resulting profile is still plausible. Extreme scenarios in which the source image does not contain enough translucency information (no noticeable shading gradients, planar surfaces with no remarkable features or strong back lighting) obviously translate into ill-conditioned linear systems that lead to erroneous profile estimations (which show as different gradients or even color shifts). Figure 7, left, shows a small object with little gradients. Conceptually, it only provides information about the leftmost part of the diffusion profile. On the other hand, Figure 7, right, presents an object illuminated from behind, which only provides info about the rightmost part of the profile. Both cases translate into numerical instability of the linear system and therefore lead to wrong captures.

Furthermore, by using the diffusion approximation, our work assumes that objects are homogeneous and optically thick, which is not the case for very small objects, or areas




Figure 7: Examples of failure cases. Left: source apple rendering input with poor subsurface scattering information and its captured material. Right: source marble rendering with strong backlight and its estimated material. The lack of information on the image or breaking our initial assumptions may lead to wrong profiles, even in controlled setups.

that present sharp edges and high curvature surfaces. Violating these assumptions may lead again to wrong profiles, or even make the QMR iterative method fail to converge.

Our method can also be potentially used in an image-editing context, by transferring the captured profile in an image object to another. By applying the same depth estimation technique both to the source and target objects, a new depiction of the latter can be created (see Figure 8). The main drawback of the technique is the double depth estimation process, which tends to accumulate larger errors in the final result. We believe this is an interesting line of future work.



Figure 8: Example of our technique as an image-editing tool. From left to right: original photograph, transfer of the wax material from the candle to the owl, and transfer from the purple wax in Figure 1 to the owl.

6. Conclusions

The approach presented in this work allows us to approximate a representation of multiple subsurface scattering in optically thick, homogeneous materials from a single image. In the absence of any prior knowledge (geometry and lighting), we face an extremely ill-posed scenario, where a physically accurate solution is simply impossible to obtain. We have shown how to overcome such scenario and still obtain good results, offering an attractive balance between visual accuracy and ease of use. Our acquired data can be directly used for rendering, while also offering a potentially interesting application as an image-editing tool.

Future research lines include the extension of our technique to heterogeneous materials or more complex BSSRDF models. Our method will also benefit from future advances in image-based light detection and depth extraction algorithms. This will allow us to extend our results to more complex objects in a wider range of scenarios. In any case, we

believe that the range of materials shown demonstrate the current practicality of the method, and hope that the contributions of this paper inspire new research in this and other related areas.

Acknowledgements

The authors would like to thank Francho Melendez for his help producing a preliminary version of the supplementary video, plus the anonymous reviewers for their valuable comments. This research has been funded by a Marie Curie grant from the Seventh Framework Programme (grant agreement no.: 251415), the Spanish Ministry of Science and Technology (TIN2007-63025 and TIN2010-21543) and the Gobierno de Aragón (projects OTRI 2009/0411 and CTPP05/09). Jose I. Echevarria was additionally funded by a research grant from the Instituto de Investigación en Ingeniería de Aragón.

References

- [BBC*94] BARRETT R., BERRY M., CHAN T. F., DEMMEL J., DONATO J., DONGARRA J., POZO R., EIJKHOUT V., VAN DER VORST H., ROMINE C.: *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, 2nd Edition. SIAM, 1994. 2, 4
- [BG01] BOIVIN S., GAGALOWICZ A.: Image-based rendering of diffuse, specular and glossy surfaces from a single image. *ACM Transactions on Graphics (SIGGRAPH)* (2001), 107–116. 2
- [dLE07] D'EON E., LUEBKE D., ENDERTON E.: Efficient rendering of human skin. In *Eurographics Symposium on Rendering* (Grenoble, France, 2007), Kautz J., Pattanaik S., (Eds.), Eurographics Association, pp. 147–157. 4
- [DRS07] DORSEY J., RUSHMEIER H., SILLION F.: *Digital Modeling of Material Appearance*. Morgan Kaufmann/Elsevier, 2007. 2
- [DS03] DACHSBACHER C., STAMMINGER M.: Translucent shadow maps. In *Proceedings of the 14th Eurographics workshop on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2003), EGRW '03, Eurographics Association, pp. 197–201. 3
- [DWd*08] DONNER C., WEYRICH T., D'EON E., RAMAMOORTHY R., RUSINKIEWICZ S.: A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM Transactions on Graphics (SIGGRAPH Asia)* 27, 5 (2008), 1–12. 2
- [FB05] FLEMING R. W., BÜLTHOFF H. H.: Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception (TAP)* 2, 3 (2005), 346–382. 5
- [GHP*08] GHOSH A., HAWKINS T., PEERS P., FREDERIKSEN S., DEBEVEC P.: Practical modeling and acquisition of layered facial reflectance. *ACM Transactions on Graphics (SIGGRAPH Asia)* 27, 5 (2008), 1–10. 2
- [GJJD09] GUTIERREZ D., JENSEN H. W., JAROSZ W., DONNER C.: Scattering. In *ACM SIGGRAPH ASIA 2009 Courses* (New York, NY, USA, 2009), SIGGRAPH ASIA '09, ACM, pp. 15:1–15:620. 1
- [GLL*04] GOESELE M., LENSCH H. P. A., LANG J., FUCHS C., SEIDEL H.-P.: Disco: acquisition of translucent objects. *ACM Transaction on Graphics (SIGGRAPH)* (2004), 835–844. 2




Figure 9: The estimated BSSRDF for the grape material in Figure 1, used to render different geometries under different lighting conditions.



Figure 10: Results of our algorithm. The small insets show the original images where the material properties are acquired from (please refer to the supplementary material for the complete data). In reading order, blue soap, whole milk, purple soap, ketchup, orange juice, whitish soap, liquid detergent, skin and greenish soap.

- [GSLM*08] GUTIERREZ D., SERON F. J., LOPEZ-MORENO J., SANCHEZ M. P., FANDOS J., REINHARD E.: Depicting procedural caustics in single images. *ACM Transaction on Graphics (SIGGRAPH Asia)* 27, 5 (2008), 1–9. [5](#)
- [JB02] JENSEN H. W., BUHLER J.: A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics (SIGGRAPH)* (2002), 576–581. [3](#)
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. *ACM Transactions on Graphics (SIGGRAPH)* (2001), 511–518. [1, 2, 4, 5, 7](#)
- [Joh02] JOHNSTON S. F.: Lumo: illumination for cel animation. In *NPAR '02: Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering* (New York, NY, USA, 2002), ACM, pp. 45–52. [6](#)
- [KRFB06] KHAN E. A., REINHARD E., FLEMING R. W., BÜLTHOFF H. H.: Image-based material editing. *ACM Transactions on Graphics (SIGGRAPH)* (2006), 654–663. [3, 5, 6](#)
- [LB00] LANGER M., BÜLTHOFF H. H.: Depth discrimination from shading under diffuse lighting. *Perception* 29, 6 (2000), 649–660. [6](#)
- [LKG*03] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics (TOG)* 22, 2 (2003), 234–257. [2](#)
- [LMHRG10] LOPEZ-MORENO J., HADAP S., REINHARD E., GUTIERREZ D.: Compositing images through light source detection. *Computers & Graphics In press* (2010). [6](#)
- [LMSSG10] LOPEZ-MORENO J., SUNDSTEDT V., SANGORRIN F., GUTIERREZ D.: Measuring the perception of light inconsistencies. In *Symposium on Applied Perception in Graphics and Visualization (APGV)* (2010), ACM Press. [6](#)
- [MMTG09] MUNOZ A., MASIA B., TOLOSA A., GUTIERREZ D.: Single-image appearance acquisition using genetic algorithms. In *Proceedings of Computer Graphics, Visualization, Computer Vision and Image Processing* (2009), CGVCVIP '09, pp. 24–32. [2](#)
- [MSY09] MUKAIGAWA Y., SUZUKI K., YAGI Y.: Analysis of subsurface scattering based on dipole approximation. *IPSJ Transactions on Computer Vision and Applications* 1 (2009), 128–138. [2, 4](#)
- [NKGR06] NAYAR S., KRISHNAN G., GROSSBERG M. D., RASKAR R.: Fast Separation of Direct and Global Components of a Scene using High Frequency Illumination. *ACM Transactions on Graphics (SIGGRAPH)* 25, 3 (2006), 935–944. [2](#)
- [PvBM*06] PEERS P., VOM BERGE K., MATUSIK W., RAMAMOORTHY R., LAWRENCE J., RUSINKIEWICZ S., DUTRÉ P.: A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics (SIGGRAPH)* (2006), 746–753. [2](#)
- [SMGG01] SLOAN P.-P. J., MARTIN W., GOOCH A., GOOCH B.: The lit sphere: a model for capturing npr shading from art. In *Graphics Interface 2001* (Toronto, Ont., Canada, Canada, 2001), GRIN'01, Canadian Information Processing Society, pp. 143–150. [2](#)
- [ST06] SHEN L., TAKEMURA H.: Spatial reflectance recovery under complex illumination from sparse images. In *Computer Vision and Pattern Recognition* (2006), IEEE, pp. 1833–1838. [2](#)
- [STPP09] SONG Y., TONG X., PELLACINI F., PEERS P.: Subedit: A representation for editing measured heterogeneous subsurface scattering. *ACM Transactions on Graphics (SIGGRAPH)* 28, 3 (2009), 1–10. [2, 6](#)
- [TGL*06] TARIQ S., GARDNER A., LLAMAS I., JONES A., DEBEVEC P., TURK G.: Efficient estimation of spatially varying subsurface scattering parameters. In *Workshop on Vision, Modeling, and Visualization (VMV)* (Aachen, Germany, 2006). [2](#)
- [WLL*08] WEYRICH T., LAWRENCE J., LENSCH H., RUSINKIEWICZ S., ZICKLER T.: Principles of appearance acquisition and representation. In *ACM SIGGRAPH 2008 classes* (New York, NY, USA, 2008), ACM, pp. 1–119. [1, 2](#)
- [WMP*06] WEYRICH T., MATUSIK W., PFISTER H., BICKEL B., DONNER C., TU C., MCANDLESS J., LEE J., NGAN A., JENSEN H. W., GROSS M.: Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics (SIGGRAPH)* 25, 3 (2006), 1013–1024. [2](#)
- [WS02] WANG Y., SAMARAS D.: Estimation of multiple illuminants from a single image of arbitrary known geometry. In *European Conference on Computer Vision (ECCV)* (2002), Springer, pp. 272–288. [6](#)
- [WT06] WU T.-P., TANG C.-K.: Separating subsurface scattering from photometric image. In *International Conference on Pattern Recognition (ICPR)* (2006), IEEE Computer Society, pp. 207–210. [2](#)
- [WZT*08] WANG J., ZHAO S., TONG X., LIN S., LIN Z., DONG Y., GUO B., SHUM H.-Y.: Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Transactions on Graphics* 27, 1 (2008), 1–18. [2, 3](#)
- [XGL*07] XU K., GAO Y., LI Y., JU T., HU S.-M.: Real-time homogenous translucent material editing. *Computer Graphics Forum* 26, 3 (2007), 545–552. [3, 4](#)
- [ZY01] ZHANG Y., YANG Y.-H.: Multiple illuminant direction detection with application to image synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 8 (2001), 915–920. [6](#)