

This is a repository copy of *Perceptually Validated Cross-Renderer Analytical BRDF Parameter Remapping*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/153581/>

Version: Accepted Version

Article:

Guarnera, Darya, Guarnera, Giuseppe Claudio, Toscani, Matteo et al. (4 more authors) (2018) Perceptually Validated Cross-Renderer Analytical BRDF Parameter Remapping. IEEE Transactions on Visualization and Computer Graphics. pp. 2258-2272. ISSN 1077-2626

<https://doi.org/10.1109/TVCG.2018.2886877>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Perceptually Validated Cross-Renderer Analytical BRDF Parameter Remapping

Dar'ya Guarnera, Giuseppe Claudio Guarnera, *Member, IEEE*, Matteo Toscani, Mashhuda Glencross, Baihua Li, Jon Yngve Hardeberg, and Karl R. Gegenfurtner,

Abstract—Material appearance of rendered objects depends on the underlying BRDF implementation used by rendering software packages. A lack of standards to exchange material parameters and data (between tools) means that artists in digital 3D prototyping and design, manually match the appearance of materials to a reference image. Since their effect on rendered output is often non-uniform and counter intuitive, selecting appropriate parameterisations for BRDF models is far from straightforward. We present a novel BRDF remapping technique, that automatically computes a mapping (BRDF Difference Probe) to match the appearance of a source material model to a target one. Through quantitative analysis, four user studies and psychometric scaling experiments, we validate our remapping framework and demonstrate that it yields a visually faithful remapping among analytical BRDFs. Most notably, our results show that even when the characteristics of the models are substantially different, such as in the case of a phenomenological model and a physically-based one, our remapped renderings are indistinguishable from the original source model.

Index Terms—BRDF model, Perceptual validation, Virtual Materials, Surface perception.

1 INTRODUCTION

In most 3D applications, the appearance of a homogeneous, opaque surface is represented by the Bidirectional Reflectance Distribution Function (BRDF) [1]. A number of BRDF models and resources are available, however they differ in the class of materials they can represent. There is no efficient general BRDF model that can be used for all classes of materials [2].

Within automotive and similar industries, design and development teams routinely use virtual 3D models which evolve through collaboration among several departments (for instance design, marketing and development). In this workflow, it is common to use many different commercial and in-house 3D tools. Traditionally, little has been done to facilitate consistent exchange of material models that preserve appearance across different renderers. This is particularly complex, since the same material model can be implemented differently between applications. Indeed, even within the same rendering tool different shader implementations might be used between versions. Consequently, there is no way to guarantee visual consistency between different applications used in a typical work-flow pipeline [3], [4].

Given a material represented using a specific (*source*) model, that is unavailable or implemented differently, a digital artist can only manually match appearance using a model available in the rendering tool in use (*target*). This involves visually matching and assigning a new set of parameters, to deliver a result that is as close

as possible to the source reference material. Since their effect on rendered images is subject to non-uniformities [5] and might also depend on scene geometry [6], selecting appropriate parameters for a BRDF model is far from straightforward. Further, due to the physical interpretation of the parameters and their scale, the effect of parameters with the same name in two different BRDFs can visually differ substantially. Thus, when an artist matches material appearance using only a visual reference encoding BRDF parameters, guessing how to replicate this is time consuming and error prone.

In this paper, we propose a novel, automatic, image-based solution to find the best set of parameters for the target material model. As input, our method uses a source BRDF model and a few High Dynamic Range (HDR) images of a rendered sphere (source BRDF probe). We define the term *BRDF Difference Probe* to encode the mapping between source and target BRDF models. Using an integer Genetic Algorithm (GA) based optimisation, we find a set of parameters in the target BRDF model, such that an object rendered using these parameters visually matches the appearance of the source. A visually indistinguishable match is only possible if the source and target BRDF models are sufficiently compatible and cannot be guaranteed in the general case. Our GA is driven by a computational metric of similarity, defined in image-space, which compares renderings of a reference scene under specially designed incident lighting.

The main components of our proposed solution are: a reference scene, with known geometry and lighting, a set of renderings (using known parameters) of the reference scene with the source material model, a GA optimisation that can access the target renderer and material model, and a fitness function driven by a computational metric, accounting for the visual differences between the source and the target model renderings. The efficacy of our approach is evaluated through numerical validation, user studies and psychometric scaling experiments. We demonstrate that renderings of our target models are visually indistinguishable from renderings using the source BRDF model (Figure 1), even

- D. Guarnera, G.C. Guarnera and J. Y., *Computer Science, Norwegian University of Science and Technology, Gjøvik, Norway, 2815.*
E-mail: {darya.guarnera, giuseppe.guarnera, jon.hardeberg}@ntnu.no
The first two authors contributed equally to this work.
- M. Toscani and K. R. Gegenfurtner, *Psychology, Justus-Liebig-Universität Giessen, Giessen, Germany, 35394.*
E-mail: {toscani, gegenfurtner}@uni-giessen.de
- M. Glencross, *Pismo Software Ltd., Oxford, United Kingdom.*
E-mail: mglencross@siggraph.org
- B. Li, *Computer Science, Loughborough University, Loughborough, United Kingdom, LE11 3TU.*
E-mail: b.li@lboro.ac.uk

Manuscript received March 24, 2018

when the parameters are spatially varying (SVBRDF).

2 RELATED WORK

Several techniques aim to facilitate material appearance design, for instance providing an intuitive painting interface like BRDF-shop [8], or through retargeting techniques that warp the reflectance properties of a source material according to the reflectance properties of a template material [9]. There is a significant body of prior work on material editing, we refer the reader to a recent survey [6].

The purpose of BRDF remapping techniques is different, since they focus on parameter specification for analytical BRDF models. Ngan *et al.* [5] defined an image-based L^2 metric to measure the distance between BRDFs. This is used to navigate variations within a model and on the manifold spanned by analytical BRDFs; reflectance neighbours in other models are found by using pre-computed conversions and multi-linear interpolation. The use of a L^2 norm metric is common in Computer Graphics for fitting measured data to a BRDF model [10], [11] and for parameter remapping [4]. Another commonly used metric, proposed by Lafortune *et al.* [12], tends to over-fit near the mirror direction and is unsuitable for applications in which human perception is important [5].

Sztrajman *et al.* [4] address the problem of remapping isotropic BRDF parameters. Their framework performs a non-linear optimisation using the trusted region reflective method. Starting from an initial guess of the parameters for the target model, they find a set of parameters that minimises the L^2 metric [10] in image space, by comparing renderings of a scene with point light illumination. We highlight the differences between Sztrajman *et al.* and our work in Section 3.6, after proving a thorough description of each component of our framework.

While BRDF fitting has some analogies with BRDF remapping [4], they are very different techniques. Both aim to find a set of parameters for a BRDF model; however fitting begins with measured data, whereas remapping takes as input rendered images. Instead of densely sampling material reflectance data, a remapping technique samples both the source and target BRDFs parameters space. In this case virtually each point represents a different material, leading to a mapping between the two spaces.

3 PROPOSED REMAPPING FRAMEWORK

Our framework is aimed at parametric BRDF models of the following form, the most common case among phenomenological and physically based models [2], [13] and widely adopted in current rendering systems:

$$\rho(\mathbf{v}_i, \mathbf{v}_r) = \rho_d + \rho_s(\mathbf{v}_i, \mathbf{v}_r) \quad (1)$$

where ρ is the overall reflectance, ρ_d and ρ_s are the diffuse and specular reflectance, \mathbf{v}_i and \mathbf{v}_r are respectively the incoming and outgoing directions. This excludes models such as Oren-Nayar [14] for diffuse, retro-reflective materials, non-parametric [15] and data driven models [2], [13]. These are outside the scope of our current remapping method.

We employ a GA optimisation, with access to a renderer implementing the target BRDF model, see Figure 2. Formally, given a distance metric $M(\cdot, \cdot) \in \mathbb{R}$ in image space, a source model S controlled by m parameters $\{p_{S1}, p_{S2}, \dots, p_{Sm}\} \in \mathbf{P}_S$, a target model T controlled by n parameters $\{p_{T1}, p_{T2}, \dots, p_{Tn}\} \in \mathbf{P}_T$, $m, n \in \mathbb{N}$, $m \neq n$ in general, remapping the source parameters into

the target ones means finding a function (BRDF Difference Probe) $f_R : \mathbf{P}_S \rightarrow \mathbf{P}_T$ such that, $\forall \{p_{S1}, p_{S2}, \dots, p_{Sm}\} \in \mathbf{P}_S$:

$$\begin{aligned} f_R(\{p_{S1}, \dots, p_{Sm}\}) &= \{p_{T1}^*, \dots, p_{Tn}^*\} = \\ &= \arg \min \{M(I_S(p_{S1}, \dots, p_{Sm}), I_T(p_{T1}, \dots, p_{Tn}))\} \end{aligned} \quad (2)$$

where I_S and I_T are renderings of the same reference scene obtained respectively with S and T .

Input to our system is simply a set of HDR renderings of a reference scene using the source BRDF model, potentially rendered using other rendering applications or other BRDF models supported by the same renderer, together with the list of parameters used for the source model. Our GA begins with a random guess for the parameters of the target model and accesses the target renderer to produce a set of images using these parameters. Driven by a fitness function, which measures the visual difference between the input and the output renderings, the optimisation requests a new set of parameters to test, and a new set of renderings is produced. This process is repeated until termination criteria, related to the average relative change in the best fitness function value over a given number of generations, are met and yields our BRDF Difference Probe.

3.1 Reference Scene Geometry, Lighting and Renders

The reference scene used by the fitness function to perform comparisons between the source and target models contains a sphere, placed on a diffuse checker-board; the incident illumination is, in turn, a point light source and an environment map.

There are a number of factors we consider with the choice of illuminant [16]. A point light source, with intensity that decreases with the square of the distance and provides a smooth falloff in all directions, allows us to assess how the BRDF spreads incoming light over the surface, and it is related to parameters such as the roughness, specularity and anisotropy. On the other hand, an environment map enables us to describe the appearance of a sphere receiving light from all directions, as found in natural environments. An environment map including step edges allows us to easily relate the edge spread visible on the surface to the roughness of the surface itself. Hence, we design and use the environment map shown in Figure 3, which displays sharp edges between the white, red, green, blue and black colour bands, as well as fine-scale colour variations in the lower half, together with rich spectral content. The colours on the environment map allow a clear analysis of the surface response to different spectral bands.

Almost every implementation of BRDF models in the required form (Eq. 1) allows independent specification of parameters related to the terms ρ_d and ρ_s , thus allowing these to be rendered independently, as in [4]. Theoretically, rendering the specular and diffuse terms independently and adding them up afterwards leads to the same result as rendering them simultaneously. However, due to renderer specific implementation choices, related to energy conservation for instance, this is not always the case, as shown in Figure 4.

To account for the aforementioned issue, we opt for rendering a diffuse only image (setting all the parameters related to the specular term to zero) and a “complete” rendering, in which both the specular and the diffuse terms have non-zero values, which constitute the input to our system. This allows us to obtain the specular term by subtracting the diffuse image from the complete one (please see supplementary material for an example), further reducing the dependency of our output on implementation choices.




Figure 1: (a) Rendering of a colour checker, using the energy conserving Ward variant. The specular reflectance (range 0.05 – 0.35) decreases from the top to the bottom row in equal steps, the isotropic roughness (range 0.01 – 0.51) decreases in equal steps from the left to the right column. (b), the corresponding remapped rendering using GGX as a target model. (c) the L^2 error plot, normalised for a better visualisation. Incident lighting from the “Uffizi Gallery” light probe [7].

Overall, for each set of source parameters, the following images are provided as input:

- A diffuse only image, rendered under point light incident illumination.
- A diffuse only image, rendered under environment lighting.
- A rendering including the diffuse and specular terms, rendered under point light illumination.
- A rendering including the diffuse and specular terms, rendered environment lighting.

The same images are rendered using the target model, upon request by our GA optimisation. These together with the source ones, provide input for the fitness function.

3.2 Fitness Function

The fitness function establishes the visual difference between input and output images, for a given set of parameters. Defined in image space, it computes a L^2 -based metric between each pair of corresponding images (*e.g.* between the source diffuse image under point light illumination and the target diffuse image under point light illumination), thus implementing the metric M in Eq. 2.

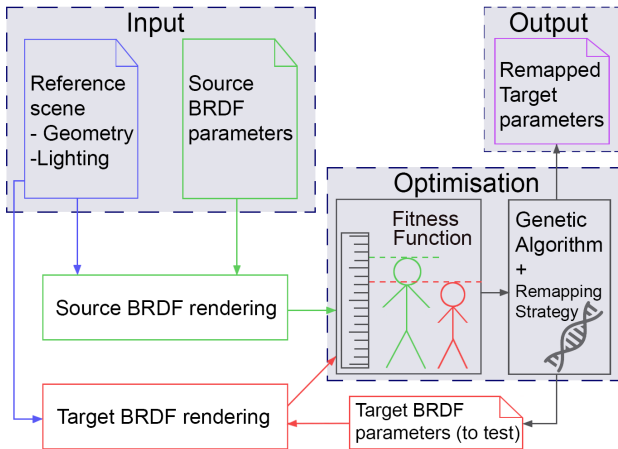


Figure 2: Architecture of our automatic BRDF parameter remapping solution.




Figure 3: Our environment map and a rendering of a grey plastic sphere.




Figure 4: In (a) rendering of a scene using both diffuse and specular terms at the same time, under environment map illumination. (b) shows the sum of the specular and diffuse terms rendered independently. In (c) the absolute difference between (a) and (b).

Table 1: Symbols used in Equation 3.

Symbol	Description
$I(\rho_d)_{SPL}$	diffuse source image, point light.
$I(\rho_d)_{TPL}$	diffuse target image, point light.
$I(\rho_d)_{SEM}$	diffuse source image, environment map.
$I(\rho_d)_{TEM}$	diffuse target image, environment map.
$I(\rho_s)_{SPL}$	specular source image, point light.
$I(\rho_s)_{TPL}$	specular target image, point light.
$I(\rho_s)_{SEM}$	specular source image, environment map.
$I(\rho_s)_{TEM}$	specular target image, environment map.
$\ \cdot\ _2$	L_2 norm.
w_D, w_S	scaling factors for all the $I(\rho_d)$ and $I(\rho_s)$ terms respectively.
α	scaling factor for the diffuse image pair, point light.
β	scaling factor for the diffuse image pair, environment map.
γ	scaling factor for the specular image pair, point light.
δ	scaling factor for the specular image pair, environment map.

For dielectric materials, the specular term tends to be, on average over the whole sphere, roughly one order of magnitude smaller than the diffuse term. Therefore, we introduce two weighting terms w_D and w_S , where $w_S > w_D$. To allow for more flexibility in the definition of the fitness function, a different weight can be assigned to each pair of corresponding terms, by introducing the additional scaling constants $\alpha, \beta, \gamma, \delta$ (see Table 1). The value f_v of the fitness function is computed using the Equation:

$$f_v = w_D (\alpha \|I(\rho_d)_{SPL} - I(\rho_d)_{TPL}\|_2 + \beta \|I(\rho_d)_{SEM} - I(\rho_d)_{TEM}\|_2) + w_S (\gamma \|I(\rho_s)_{SPL} - I(\rho_s)_{TPL}\|_2 + \delta \|I(\rho_s)_{SEM} - I(\rho_s)_{TEM}\|_2) \quad (3)$$

Table 1 contains a glossary of symbols used; Figure 5 illustrates how the computations are performed by the fitness function.

When $w_D = w_S = 1$, $\alpha = \beta = \gamma = \delta = 1$ and the diffuse and specular terms are converted into the CIELAB colour space before computing the norm, Eq. 3 is equivalent to the ΔE_{76} colour

difference formula [17].

Since Eq. 2 can be seen as a maximisation of the image fidelity between source and target renderings, M (and hence the fitness function which implements it) acts as a full-reference image fidelity measure [18]. We exploit this concept, and the flexibility allowed by the scaling factors in Eq. 3, tuning them in order to maximise the correlation between Eq. 3 and the Gradient Similarity (GM) metric [19]. See Section 4 for further details. Other than colour differences, this effectively captures local image structures, to which the human visual system is highly sensitive.

Many empirical BRDF models, such as Phong [20], Ward [21] and Lafortune *et al.* [12], do not account for the Fresnel effect, which has a significant impact on the appearance of a dielectric, especially at grazing angles (see Figure 6). This implies that a remapping from a model without the Fresnel term to a model accounting for it (and vice-versa), will lead to high errors towards object silhouettes, which might dominate the value of the fitness function and lead to inaccurate results. To obtain a good fit this must be taken into account, through an extra weighting function which gives more importance to the central area of the sphere in the reference scene, rather than the edges. We implement a weighting function as a mask (“Optional mask” in Figure 5), derived with the equation $w_p = 1 - (\|p - C\|_2/r)^{\frac{1}{2}}$, where p are the coordinates of a pixel on the sphere in the image lattice, C are the coordinates of the centre of the sphere and r is its radius.

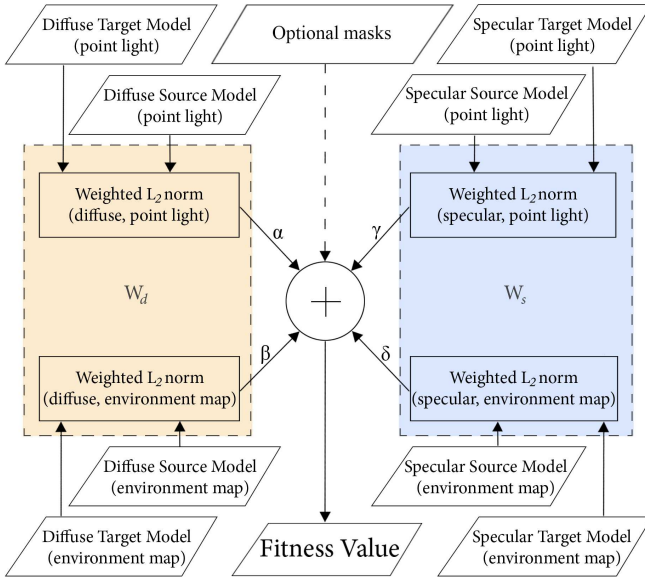


Figure 5: Computations performed by the fitness function.

3.3 Integer Genetic Algorithm Optimisation

Genetic algorithms are a class of stochastic search strategies modelled after evolutionary mechanisms. GAs have been successfully used in Computer Graphics, for instance in the context of learning new analytic BRDF models [11].

Given a fitness function to be optimised, a genetic algorithm randomly selects some points in the definition domain, which constitute the initial population. At each step of the optimisation, the genetic algorithm accounts for the fitness function and selects some individuals from the current population to be parents. From these points the next generation of children is produced, either by crossing over the vector entries of two parents (their “DNA”) or by




Figure 6: In (a), rendering of a green plastic sphere under uniform incident spherical lighting, accounting for the Fresnel term (GGX model, on the left) and without accounting for it (Ward model, on the right). In (b), the normalised specular reflectance profiles of the two models behaves in the opposite way.

applying small mutations to the “DNA” of a single parent. Finally, the best performing individuals “survive” to the next generation.

GAs offer a suitable strategy to optimise both constrained and unconstrained non-linear systems with a large number of variables, in particular when it is challenging to accurately model the interaction between them and incorporate such information into an analytical cost function [22]. Hence, GAs can be used when no information is available about the gradient of the function, which does not need to be continuous nor differentiable. They also provide good results even when the function has several local minima. The evaluation of a BRDF model potentially involves a high number of variables and the exact implementation is rarely available. Therefore, it is impossible to accurately model the interaction among all variables.

As with most optimisation techniques, a standard GA expects as input a continuous variable. However, in image-based remapping techniques, there is a high computational overhead for each point to sample in a continuous space, due to the need to render one or more images to compute the objective function [4].

Given the above, to reduce computational overhead it is desirable to set a step for each dimension involved in the optimisation, thus fixing the resolution of a regular lattice in the target BRDF parameter space. In our framework we employ integer GA optimisation, where the rules to create the initial population, parents cross-over and mutation, force the variables to be integers [23]. Given a parameter x which assumes values $x \in [\psi - \Psi, \Psi]$, a sampling resolution ζ , a mapping $\bar{x} \in \{\psi, \psi + \zeta, \dots, \Psi\} \rightarrow \mathbb{N}^+$ can be derived through $i = \lfloor (\bar{x} - \psi) / \zeta \rfloor + 1$, $i \in \{1, \dots, \lfloor (\Psi - \psi) / \zeta \rfloor\}$. Please note that other mappings from the discretised parameter domain to \mathbb{N}^+ , allow a non linear sampling. The fitness function is augmented by a parameter-free penalty term, related to the feasibility of the solution (*i.e.* whether each integer variable v_k to optimise is in the range $v_k \in \{1, \dots, \lfloor (\Psi_k - \psi_k) / \zeta_k \rfloor\}$). If a member is feasible, the penalty function is the fitness function itself; otherwise, the penalty function is the maximum fitness function among feasible members of the population, plus a sum of the constraint violations of the point. To select individuals for the next generation, pair-wise comparison in tournament selection is used. In the tournament selection, two individuals are randomly picked from the population and compared. Any feasible individual is preferred to any infeasible one; between two feasible individuals, the one with a lower value of the fitness functions is preferred. Finally, between two infeasible solutions, the one with a smaller constraint violation is preferred. Overall, the tournament selection directs the optimisation towards the feasible region [24].

3.4 Remapping Strategies

Finding the BRDF Difference Probe mapping the set of m source parameters to the set of n target parameters is achieved in two different ways, unconstrained and constrained remapping.

- (i) *Unconstrained remapping*: given a m -tuple of source parameters, allow the GA to explore any n -tuple of target parameters searching for the best mapping according to the fitness function.
- (ii) *Constrained remapping*: given a set of m -tuples of source parameters, constrain the GA to explore variations in only one source parameter at a time. The sequence of source parameters must be provided in a monotonically increasing fashion. The starting point of the search in the target parameter space makes use of previously remapped parameters. It follows that the mapping between corresponding parameters is monotonically non-decreasing: if for a given source parameter p_{S_i} , related to a target parameter p_{T_j} , a mapping $\{p_{S_i} = \kappa\} \rightarrow \{p_{T_j} = \nu\}$ is found, then for a value of $p_{S_i} = \kappa + \delta$, $\delta > 0$, a suitable mapping can only be greater or equal to $p_{T_j} = \nu$.

Strategy (i) can be immediately applied and does not require any particular care. However, (ii) involves some additional considerations and prior knowledge of the semantic relationship among parameters, but with the benefit of a progressive reduction of the search space and computational time and the inverse mapping from target to source models.

3.5 BRDF Difference Probe

Once the GA discovers a set of parameters for the target BRDF, for a given set of source renderings and parameters, these are stored in a look-up table indexed by the parameters of the source BRDF. The algorithm then starts the optimisation with a new set of source renderings and parameters. When all the source parameters/renderings provided as input have been remapped to the target BRDF, the fitting module takes care of finding a smooth, analytical relationship between the input and output parameters, through local regression smoothing. When such a relationship cannot be found, the output is the aforementioned look-up table.

3.6 Differences with Previous Work

In this Section we provide a summary of the main differences between our work and the closest related one, by Sztrajman *et al.* [4]. The first important difference is the objective function used. While in [4] a standard L^2 metric in image space is used, in our work we modify the standard definition to allow our fitness function to account (at the same time) for colour differences (ΔE_{76} [17]) and gradient differences (GSM [19]). Consequently, we exploit two features to which the HVS is extremely sensitive, thus providing a perceptually accurate remapping, as demonstrated in the remainder of the paper. Finally, our fitness function, by means of the optional masks, allows us to more robustly deal with a remapping from a model with a Fresnel term to a model without it, and vice-versa.

The second important difference is the type of optimisation used. By means of integer GA optimisation, we are able to force the search space in the target BRDF model to be a regular N-dimensional lattice, sampled at a arbitrary resolution, thus enormously reducing the number of rendered images during the optimisation. Further, we apply several computational optimisations to

reduce the computational complexity, as highlighted in Section 4, such as the use of a Principal Component Analysis (PCA) based synthesis, to increase the sampling resolution without the need for actually rendering the images.

While we demonstrate our framework on both anisotropic BRDFs and SVBRDFs, Sztrajman *et al.* [4] limit theirs to homogeneous, isotropic BRDFs. Furthermore, unlike Sztrajman *et al.*, our framework relaxes the need for users to provide an initial guess of the target parameters, which for BRDF models with a high number of parameters proves to be difficult, given also the non-uniformity of their effects [5]. We provide the user with a tool that can be used either as a “black-box” (*unconstrained remapping*) or incorporates previous knowledge about the semantic and functional relationship among source and target parameters (*constrained remapping*).

4 EXPERIMENTAL SETUP

Our framework is implemented using Matlab. Although any renderer with a scripting interface can be used for the experiments reported in this paper, we use the Mitsuba renderer [25]. All experiments are performed on a mobile PC with an Intel® Core™ i7-6700HQ CPU and 32GB memory. The images used to remap the source into the target models (BRDF Difference Probe) are rendered at 300×300 pixel resolution.

The scaling factors in Eq. 3 are computed by maximising the correlation between Eq. 3 itself and the GM metric [19] on a set of renderings used as a benchmark. We prefer to use our fitness function rather than GM since the latter is defined on low dynamic range (LDR) images, and its computation is more than 3 times slower than the fitness function (GM takes about 0.23 seconds for a pair of 512×512 LDR images, compared with 0.07 seconds required by our fitness function for two HDR images at the same resolution). Hence, we set $W_d = 0.25$, $W_s = 1 - W_d$; $\alpha = \gamma = 0.025$, $\beta = 1 - \alpha$, $\delta = 1 - \gamma$, which gives an average non-linear correlation (measured as the distance correlation [26]) of 0.783, and up to 0.943 over the benchmark.

In the following we describe optimisations designed to further reduce the computational cost of our framework.

Both Strategies. Since the GA involves some random steps, for a given source input the fitness function could potentially be computed several times for the same set of target parameters. To avoid redundant computations we implement a lookup table, indexed by the values of the source and target parameters. Every time a new set of source input parameters and renderings are provided to the remapping pipeline, if a new value of the fitness function is computed it is stored for fast access. We implement a database to keep track of all the renderings produced and the parameters used for each of them. When requested by the GA if the image already exists in the database, it is simply retrieved and provided to the fitness function.

Constrained Strategy. For both the lighting conditions selected for the reference scene, we pre-render a set of target images at a half sampling rate in each target parameter direction with respect to the rate used in the experiments. PCA is subsequently applied on the set of pre-rendered images. Since the scene geometry is fixed and the only variations are due to non-linear parameters, the first 10 principal components suffice to faithfully reconstruct the input renderings. For each rendering, the loadings derived by the PCA are fitted to its parameters, and used to synthesise the missing renderings to achieve the full sampling rate requested. This allows us to dramatically reduce the rendering

time, without affecting the remapping results. PCA is commonly used in this context [5].

5 REMAPPING OF ISOTROPIC BRDF MODELS

In this Section, we present a set of experiments involving isotropic BRDF models. We demonstrate the usefulness of our framework on both a $1D \rightarrow 1D$ remapping case, and on a more general $2D \rightarrow 3D$ case, where the dimensionality of the remapping is given by the number of parameters involved in the source and target models specular lobes; please note that the remapping also involves the diffuse component. We use this simpler cases to demonstrate the main concepts, before addressing higher dimensional cases and anisotropic models.

5.1 Isotropic BRDF Models - 1D Remapping

To test our framework on the simplest case we focus on metals, and select as a source BRDF the phenomenological Ashikhmin-Shirley model [27], whereas the target BRDF is the physically-based Cook-Torrance model [28], with a Beckmann distribution of microfacets. The user selectable parameters of the Ashikhmin-Shirley source model used for the experiments are the real and imaginary components of the material IOR (material dependent), and the anisotropic roughness along the tangent and bitangent directions (respectively $R_{U_{as}}, R_{V_{as}} \in [0 - 1]$). Since the target model (Cook-Torrance) does not support anisotropic materials, we focus on isotropic conductors, we set $R_{as} = R_{U_{as}} = R_{V_{as}}$. As for the parameters of the Cook-Torrance model, the parameters are the real and imaginary components of the material IOR (material dependent), and the isotropic roughness of the microgeometry ($R_{ct} \in [0 - 1]$).

To perform remapping we use measured real world conductors, hence the real and imaginary components of the IOR are determined by the material considered (iridium was used to find the mapping). To summarise, given the parameters fixed by the conductor selection, we search for a mapping between $\{R_{as}\} \rightarrow \{R_{ct}\}$. Given the nature of the parameters involved, the constrained strategy can be readily implemented. The parameters available for the GGX and Cook-Torrance models are the same, and we look for a mapping $\{R_{as}\} \rightarrow \{R_{ggx}\}$.

Even the simple 1D remapping allows us to highlight some challenging situations, for instance when remapping a microfacets source BRDF, based on a Gaussian distribution, to the target model GGX, which implements the Trowbridge-Reitz distribution [29].

Table 2: Min, max, mean and standard deviation σ of the NRMSD, per colour channel, between source (Ashikhmin-Shirley model, A-S) and target (Cook-Torrance, on the left side, and GGX, on the right side) renderings.

	A-S \rightarrow Cook-Torrance			A-S \rightarrow GGX		
	Red	Green	Blue	Red	Green	Blue
min	0.0054	0.0046	0.0043	0.0166	0.0154	0.0123
max	0.0305	0.0320	0.0329	0.0546	0.0483	0.0404
mean	0.0107	0.0103	0.0100	0.0247	0.0259	0.0278
σ	0.0078	0.0084	0.0087	0.0097	0.0076	0.0059

Results - Isotropic BRDF Models (1D Remapping)

The test scene used in this set of experiments consists of three spheres with the same radius, placed on a checker-board and illuminated by the ‘‘At the Window’’ [30] light probe (©Bernhard Vogel), which displays rich content (windows, statues, furniture,




Figure 7: Remapping of isotropic BRDF models, 1D case. We show the source renderings (Ashikhmin-Shirley, first column), the corresponding remapped target images for the Cook-Torrance (second column) and GGX models (third column). In the insets we report the L^2 error maps in RGB space between the source and target renderings.

etc.); a point light is also included in the scene. We assigned three different metals to the spheres in the scene (gold, silver and copper). Figure 7 shows some of the source and our remapped renderings (remapping from Ashikhmin-Shirley to Cook-Torrance), respectively on the first and second columns. Table 2 gives the per-channel Normalised Root Mean Square Deviation (NRMSD) among source and remapped renderings, computed by normalising the RMSD by the range observed in the source rendering. For reference, the average NRMSD over ten renderings of the same scene, due to the stochastic nature of the path tracer and sample generator, is 0.015.

We repeated the experiment using the GGX model as a target; while our framework managed to provide a visually satisfactory remapping from the source to the target parameters (see Figure 7, third column), the average NRMSD is up to 5.5 times higher than in the previous case and about 2.5 times on average, as reported in Table 2. This is due to the very different characteristics of the microfacets distribution at the core of the GGX model with respect to a Gaussian distribution, which is engineered to have a narrow specular peak and a much longer tail than usual. As a consequence, a rendering using the GGX model displays more shadowing than many other BRDF models [2]. Figure 8 shows three renderings of an iridium sphere, from left to right using the Ashikhmin-Shirley model, the corresponding rendering with the remapped parameter for the Cook-Torrance and the GGX model.

We note that this simple 1D case could be solved by pre-rendering the reference scenes by spanning the range of the roughness in the model (brute-force approach), and then applying a simple regression on the minima given by Eq. 3. However, this would not lead to any advantage with respect to our framework. Instead, since it would necessitate rendering all the aforementioned images, even using a constrained strategy, it is equivalent to the worst case scenario in our framework.

5.2 Isotropic BRDF Models - 2D to 3D Remapping

To test the robustness of the proposed framework on the $2D \rightarrow 3D$ remapping of isotropic BRDF models, we present as a case study the remapping from the Ward BRDF [21], in particular the Geisler-Moroder-Dür energy conserving variant [31] and the BRDF component of the GGX model [32], focusing on dielectric materials.




Figure 8: Visual comparison about the effect of comparable values of surface roughness on the Ashikhmin-Shirley (A-S), Cook-Torrance and GGX models.

Table 3: Parameters available to the user, in the implementation of the Ward and GGX models used in the isotropic 2D to 3D remapping experiments.

Ward parameters (dielectrics)		
Parameter description	Symbol	Plausible values
Diffuse reflectance	D_w	$[0 - 1]$
Specular reflectance	S_w	$[0 - 0.4]$
Anisotropic roughness along the tangent direction	R_{U_w}	$]0 - 0.7]$
Anisotropic roughness along the bitangent direction	R_{V_w}	$]0 - 0.7]$
GGX parameters (dielectrics)		
Parameter description	Symbol	Plausible values
Diffuse reflectance	D_{ggx}	$[0 - 1]$
Specular reflectance	S_{ggx}	0, 1
Index of Refraction of the material	I_{ggx}	$]1 - 1.85]$
Isotropic roughness of the surface microgeometry	R_{ggx}	$[0 - 0.7]$

We select as source and target BRDFs two very different models, the first one being a phenomenological model, not based on the microfacets theory and with no support for the Fresnel effect. The second one is a physically based model which describes a microfacets distribution and accounts for the Fresnel effect. We believe this case study is of particular interest not only in computer graphics, but also in the colorimetry and human vision communities for instance, where phenomenological models and in particular the Ward BRDF, are widely used [33], [34].

Parameters which can be set by the user, in the implementation of the Ward and GGX models used for this experiment are listed in Table 3. Plausible values for dielectrics are derived from the Mitsuba manual and empirically through additional experiments. Since the implementation of the GGX model used for this experiment does not support anisotropic materials, we use the isotropic version of the Ward model, by setting the roughness along the tangent and bitangent directions to the same value. To derive the mapping in both models the diffuse reflectance colour is fixed to a neutral value (grey), hence we aim to find a mapping between the following sets of parameters:

$$P_S \rightarrow P_T; P_S = \{S_w, R_w\}, P_T = \{S_{ggx}, I_{ggx}, R_{ggx}\}, R_w = R_{U_w} = R_{V_w}.$$

The resolutions for source and target parameters used in our experiments, spanning the range of plausible values (Table 3), is reported in the supplementary material. No additional information is required for the unconstrained remapping strategy; to implement the constrained one we observe that:

- The specular reflectance in the GGX model, to fulfil energy conservation, can be either 0 (lambertian material) or 1 (with the Index of Refraction controlling the specularity). In the Ward model this is equivalent to setting the specular

- reflectance respectively to null or any non-null positive value.
- A non-null specular reflectance in the Ward model can only be related to the Index of Refraction (IOR) in the GGX model, which controls the Fresnel effect and hence the amount of energy reflected through pure specular reflection.
- The isotropic roughness in the Ward model can only be related to the microgeometry roughness in the GGX model.

Results - Isotropic BRDF Models (2D to 3D Remapping)

The test scene used for this set of experiments is the same as in the 1D case; each sphere in the scene has a different diffuse reflectance (Figure 9). In the first and second columns of Figure 9, we report a visual assessment of the quality of the unconstrained remapping, showing some of the source images, the corresponding remapped target images and the L^2 error maps in RGB space (in the insets); Table 4, top half, summarises NRMSD for all the images used in the first experiment.




Figure 9: Remapping of isotropic Ward to GGX. We show the source renderings (first column), the corresponding remapped target images for the GGX model, using the unconstrained setting (second column) and the constrained one (third column). The first row refers to a source roughness of 0.31, whereas 0.01 is used in the second row. In the insets we report the L^2 error maps.

Table 4: Quantitative validation of the remapping strategies - isotropic dielectric materials. Min, max, mean and standard deviation σ of the NRMSD, per colour channel, between source and target renderings.

	Unconstrained remapping			Constrained remapping		
	Red	Green	Blue	Red	Green	Blue
min	0.0105	0.0133	0.0114	0.0113	0.0144	0.0114
max	0.0323	0.0374	0.0440	0.0306	0.0366	0.0440
mean	0.0182	0.0222	0.0262	0.0187	0.0230	0.0268
σ	0.0059	0.0068	0.0097	0.0056	0.0067	0.0102

For the constrained remapping experiment, Figure 9 shows a visual assessment of the quality of the remapping (compare the first and third columns); Table 4, bottom half, summarises NRMSD for all the images used in the experiment.

In both experiments most of the error is localised around the edges of the spheres, as shown in the insets in Figure 9. This is due to the lack of support for the Fresnel effect in the Ward model and, as previously explained, cannot be fixed by a remapping algorithm.

Figure 10 shows an example of remapping a spatially varying BRDF (SVBRDF), in which the roughness is controlled by a texture map; once this remapping has been performed as described in the above, its output can be applied on per texel basis. Similar considerations apply to the spatially varying specularity (Fig. 1).




Figure 10: Remapping of SVBRDFs. The roughness in the source rendering (a) is controlled by a texture map. The remapped rendering is obtained by applying the mapping to each texel of the roughness map.

Besides a similar NRMSD (Table 4), an important advantage of the constrained strategy is the ability to derive smooth relationships between corresponding parameters, whereas the unconstrained remapping strategy might force the fitting module to output a look-up table, since all the target parameters are allowed to vary at once. This is particularly noticeable for the selection of diffuse only/diffuse and specular (Figure 11(a)) and roughness (Figure 11(c)) in the remapped GGX parameters. Also the remapped IOR (Figure 11(b)) displays dependencies on both the S_w and R_w . As a further consequence, the unconstrained strategy must sample the whole parameter space of the source model, since at any stage of the remapping it is not possible to predict how a given set of input parameters will be remapped.

The constrained strategy conveys some important generalisation characteristics, which allow us to sample only a relatively small portion of the source model parameter space, as reported in Figures 11(e) and 11(f). The output of the GA is then used by the fitting module to derive the relationships among source and target parameters. Overall, this information can be inverted and used for the inverse remapping, *i.e.* $GGX \rightarrow Ward$.

In Table 5 we compare the runtime of a brute force approach, with the unconstrained and constrained strategies runtime in the worst-case scenario. In all cases we refer to a single-thread implementation; the runtime is dominated by rendering time, with the constrained strategy being the fastest and requiring approximately 9 hours to perform a full remapping, making it suitable for use as a overnight tool even on a mobile PC system. Overall, on a 48-core Intel Xeon E7-8894 v4 platform, an optimised implementation of the unconstrained remapping takes less than 1.7 hours, whereas the constrained remapping requires about 52 minutes.

5.3 Isotropic BRDF Models - Higher Dimensional Cases

In this Section we describe the Remapping from the isotropic microfacet BRDF model by Löw *et al.* [35] to GGX. While the

Table 5: Comparison of the running time for a full remapping Ward \rightarrow GGX (isotropic case), on the mobile PC described in Section 4.

	Brute Force	Unconstrained	Constrained
remapping	22h 36m	1h 28m	0h 46 mi
rendering	31h 42m	31h 42m	8h 18m
total	54h 18m	33h 10m	9h 04m

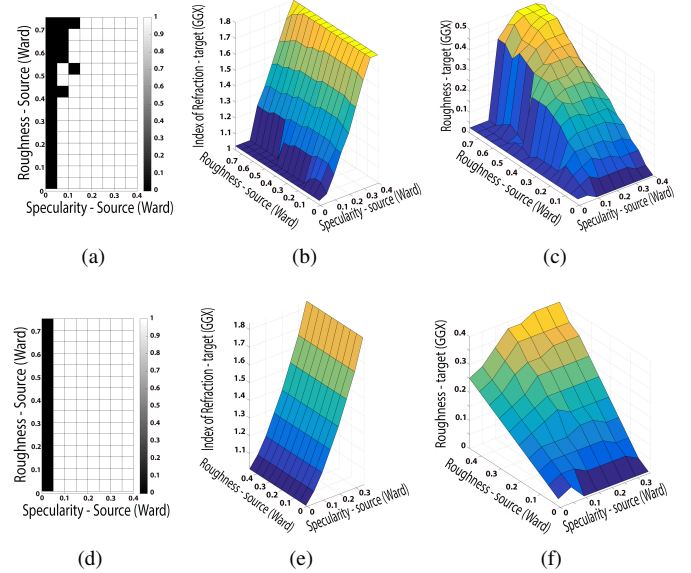


Figure 11: Comparison between the output of the unconstrained and constrained strategies. In the first row we report the output of the unconstrained remapping, showing (a) the binary choice between diffuse only (black) or diffuse and specular (white) for the remapped target parameter S_{ggx} , the remapped target I_{ggx} (b) and target R_{ggx} (c). The second row, in (c)-(f) reports the same information as derived by the constrained remapping.

Fresnel and shadowing/masking terms are the same as in Cook-Torrance [28], the microfacet distribution is based on a modified ABC distribution [36].

Other than the index of refraction $I_{L\ddot{o}w}$, The parameters that control the model are $A_{R|G|B}L\ddot{o}w$, a per-channel specular scaling term (hence, the distribution is not normalised and the model does not obey to the energy conservation principle), $B_{L\ddot{o}w}$ which controls the sharpness of the specular lobe and $C_{L\ddot{o}w}$, which controls the shape of the tail of the distribution and hence is related to the fall-off rate at wide-angle scattering. As for the GGX model, we modify it in order to allow a different $S_{R|G|B}ggx$ per channel, in the range $[0 - 1]$. Hence, a naive implementation of the remapping would lead to a $6D \rightarrow 5D$ case. However, the mapping can be derived for a single channel, and applied to other channels independently, thus reducing the dimensionality to $4D \rightarrow 3D$.

To implement the constrained strategy we observe that the $A_{R|G|B}L\ddot{o}w$ are clearly related to $S_{R|G|B}ggx$, as well as $I_{L\ddot{o}w}$ and I_{ggx} ; however, as reported by Löw *et al.* [35], $I_{L\ddot{o}w}$ does not have a physical meaning and the range of values it assumes does not correspond to real world materials. As for the roughness, the combined effect of $B_{L\ddot{o}w}$ and $C_{L\ddot{o}w}$ can be related to R_{ggx} .

Results - Isotropic BRDF Models (Higher Dimensional Remapping)

To validate the remapping from Löw *et al.*'s BRDF model (referred to hereafter as ABC) to GGX we make use of the MERL material database [37], through the following steps:

- We use the technique described in Löw *et al.* [35] to fit the MERL materials, thus finding the optimal parameters for the

ABC model to represent them;

- We apply our remapping $ABC \rightarrow GGX$ to find a suitable representation of the materials in the GGX parameters space.

As a further comparison, we show the fitting of the MERL materials to the GGX model, obtained using the method by Dupuy *et al.* [38]. To display our results, we take advantage of the perceptual validation of the optimised surface and light direction suggested by [39] (see Figure 12). Please note that our input is not a measured material. Hence, images are rendered using the surface geometry, rather than mapping measured values directly on it.

While for most MERL materials the remapping leads to convincing results, with the search space in the I_{ggx} in the range $[1.01 - 2.1]$, in 29 cases (including many metals and the specular-phenolic family) we had to extend it. This is due to the parameters in the ABC source renderings, with values of $I_{Löw}$ as high as 99.99 and $C_{Löw}$ values largely outside the typical range $[0.5 - 1.5]$ [35]. Overall, the remapping failed on the 3 gold-metallic paints.

6 REMAPPING OF ANISOTROPIC BRDF MODELS

In this Section we focus on anisotropic BRDF models and show remapping between Lafortune to Ward and GGX to Ward models.

6.1 Anisotropic BRDF models - Lafortune to Ward

The empirical BRDF model by Lafortune *et al.* [12] is a generalisation of the cosine lobe model, allowing multiple steerable lobes. Other than the specular reflectance S_L , each of the specular lobes l is controlled by the parameters $C_{x,l}$, $C_{y,l}$, $C_{z,l}$ and the exponent n_l , which controls the shape of the lobe. The $C_{x,l}$, $C_{y,l}$, $C_{z,l}$ parameters allow control of retro-reflections (by setting them all to positive values), anisotropy (when $C_{x,l} \neq C_{y,l}$) and off-specular peaks ($C_{z,l} \leq -C_{x,l} = -C_{y,l}$); the original Phong model [20] can be obtained by setting $\|l\| = 1$, $C_x = C_y = -1$, $C_z = 1$, whereas the Lambertian model can be obtained by setting $n_l = 0$.

We demonstrate the remapping from the Lafortune BRDF model to Ward, in the anisotropic case, single lobe. For the specular lobe we look for a mapping from $\{S_L, C_x, C_y, C_z, n\} \rightarrow \{S_w, R_{U_w}, R_{V_w}\}$. Since the anisotropic Ward model does not account for retro-reflection and off-specular peaks, we exclude the corresponding Lafortune parameters from our analysis. For the constrained remapping strategy, S_L and S_w are corresponding parameters. The anisotropy in Lafortune is controlled by C_x and C_y , with n controlling the roughness, whereas in Ward anisotropy and roughness are controlled concurrently by R_{U_w} and R_{V_w} .

Results - Anisotropic Lafortune to Ward

In Figure 13, the top row, shows a nuts and bolts cat rendered using the Lafortune model, for different degrees and direction of anisotropy. The anisotropy in the tangent directions decreases from the leftmost image towards the centre one, which is isotropic, then the anisotropy in the bitangent direction progressively increases moving towards the rightmost image. The second row shows the same images rendered with the anisotropic Ward model, where the parameters have been derived by our remapping framework.

While our remapping is able to produce visually satisfactory results for many points in the Lafortune parameter space, the decoupling between roughness and anisotropy leads to failure when trying to remap glossy, anisotropic surfaces (*i.e.* $n > 128$, $C_x \neq C_y$) to the Ward model. In the latter, roughness in the tangent/bitangent directions and anisotropy are strictly correlated. No issues arise when remapping isotropic surfaces.

6.2 Anisotropic BRDF models - Ward to GGX

We present the remapping from the anisotropic Ward model to the anisotropic GGX [40]. The remapping is performed between the two sets of parameters $\{S_w, R_{U_w}, R_{V_w}\} \rightarrow \{S_{ggx}, R_{U_{ggx}}, R_{V_{ggx}}\}$ (see Table 3), where $R_{U_{ggx}}$ and $R_{V_{ggx}}$ extend R_{ggx} to the anisotropic case. Given the implementation of GGX differs from the one used in Section 5.2, previous remapping results cannot even partially be reused. As in the corresponding isotropic case, we make use of the additional mask in the fitness function computation, to reduce the effect of the Fresnel in the remapping.

Results - Anisotropic Ward to GGX

The second row of Figure 13, shows a nuts and bolts cat rendered using the Ward model, as explained in the previous section. The third row shows the same images rendered with the anisotropic GGX model, where the parameters have been derived by our remapping framework. Besides the unavoidable differences due to the Fresnel effect, our framework provides satisfactory results.

7 PERCEPTUAL VALIDATION: USER STUDIES AND PSYCHOMETRIC SCALING EXPERIMENTS

A quantitative validation provides an initial assessment of the quality of the remapping performed by our framework. However, to ensure that the remapping produces visually plausible results, regardless of the selected metric, results are also verified through user studies and psychometric scaling experiments, focusing on the isotropic BRDF models remapping.

7.1 User Studies

User studies are a common way to validate results in computer graphics, with a range of available study methodologies and metrics [41], we use similarity judgements. The goal of our user studies is to establish if participants can distinguish between the appearance of an image rendered using the source BRDF and the same image rendered using the target BRDF, where the parameters are derived by our remapping framework. In our study, participants with corrected to normal vision are presented with a sequence of two images next to each other (stimuli) on a uniform grey background, on a calibrated display.

Each image pair consists of a scene rendered using the source BRDF model and the same scene rendered with the target BRDF using our remapped parameters. Each participant, after observing an image pair for exactly 5 seconds, was asked to rate the similarity between displayed images, using integer ratings on a scale 0 – 2, where 0 means “noticeably different”, 1 encodes “slightly different” and 2 means “same”. Instructions were provided before the start of the study. Image pairs were shown to subjects in a randomised order and included a control set of image pairs in which the images are actually the same, to provide a reference for statistical analysis. Four user studies were conducted:

- (1) Unconstrained remapping, dielectrics. The study used 36 test pairs, generated by the unconstrained remapping experiment in Section 5.3 (Figure 14(a)), and 5 control pairs.
- (2) Constrained remapping, dielectrics. Source and target BRDFs were the same as in the previous user study, but with different diffuse colours for the spheres (Figure 14(b)). Like the first study, participants were shown 41 pairs (36 test, 5 control).
- (3) The third user study presents a set of conductors (metals) with varying roughness, where the remapping from the source




Figure 12: The first column shows the measured MERL materials [37] applied to the scene with optimised geometry and lighting by Havran *et al.* [39]. The second column reports the same scene rendered using the ABC model, where the model parameters have been derived by fitting as described in Löw *et al.* [35]; these renderings represent the source for our remapping. The third column reports the scene rendered using the GGX model, using the parameters derived from our remapping ABC → GGX. As a further comparison, the last column reports the scene rendered using the GGX model, with the parameters derived from the fitting described in Dupuy *et al.* [38]. Please see the supplementary material for the complete list of materials in the MERL dataset.



Figure 13: A nuts and bolts cat rendered using the Lafortune model, for different degrees and direction of anisotropy. The second row shows the same images rendered with the anisotropic Ward model, whereas the bottom row show the images rendered using the anisotropic GGX implementation.

(Ashikhmin-Shirley) to the target model (Cook-Torrance) is performed with the constrained remapping strategy (Figure 14(c)). In this case the mapping is performed only in one dimension ($R_{as} \rightarrow R_{ct}$); 19 test pairs generated by the

unconstrained remapping experiment in Section 5.1 were used together with 2 control pairs.

- (4) In the final user study the shapes placed in the scene are the “Stanford Bunny” [42], “Happy Buddha” [43] and “Dragon”

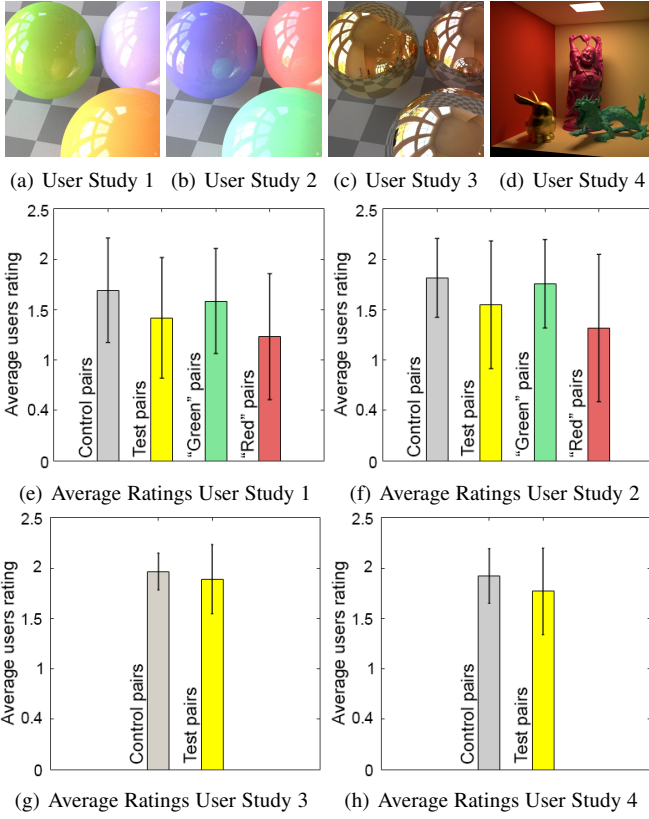


Figure 14: (a)-(d) Examples of the test scenes used in the user studies; (e)-(h) for each study we report the mean rating and the standard deviation. The grey bar refers to the control pairs (same-image pairs), the yellow bar refers to the ratings for all the test pairs (original-remapped). For studies 1 and 2 we give the mean and standard deviation of the “green” and “red” groups of test pairs.

(Figure 14(d)). For each test pair, the source image was generated by randomly assigning to the objects one of the source BRDF models Ward or Ashikhmin-Shirley, with parameters randomly selected as well; the target image was generated by assigning to each object the corresponding target BRDF model and remapped parameters. Participants were presented with 2 test pairs and 2 control pairs.

Participants reported that, for most image pairs, they had difficulties spotting differences between the left and right renderings. In fact, in many cases, similar average scores were given to the original-remapped image pairs and to the control pairs.

Results of these user studies were analysed for statistical significance, using both the one-tailed t -test and Welch’s t -test, accounting for both the possibilities that the variances of the test and control pairs are equal or unequal. Additionally, for each user study we compared ratings given by users to the test pairs and to the control pairs using the Wilcoxon signed rank test. All tests were conducted at a significance level $\alpha = 0.05$. A detailed report of all the statistical tests is included in supplementary material, we report the main findings below.

User Studies 1 and 2. The pairwise comparison results from a multiple comparison test using the Tukey’s honestly significant difference procedure, conducting the hypothesis tests at the 5% significance level, allowed us to highlight that participants tend to spot differences between source and remapped image pairs when $0 < R_w \leq 0.06$, or when $S_w \geq 0.35$. Accordingly, we partitioned the test image pairs into *red* = $\{\{0 \leq R_w \leq 0.06\} \cap \{S_w \geq 0.35\}\}$ and *green* = $\{\{R_w > 0.06\} \cap \{S_w < 0.35\}\}$ groups. Both the t -test

and Welch’s t -test confirm that for renderings belonging to the green area (a large portion of the parameter space) observers are unable to reliably distinguish between source and remapped renderings, irrespective of the remapping strategy. As for the red group, both tests reject the null hypothesis; our results are further confirmed by the Wilcoxon test. For details about the partitioning of the Ward parameter space and the statistical tests, see supplementary material.

User Study 3. The average rating for all the images is reported in Figure 14(g). The statistical tests, performed by including all the test pairs in one group, indicated no significant difference among test and control pairs, confirming participants’ impressions that most image pairs appeared to show the same image.

User Study 4. The purpose of the fourth user study was to assess the efficacy of our framework when combining several BRDF models and complex objects in one scene. Both tests, performed by including all the test pairs in one group, indicated no significant difference among the rating of test and control pairs.

7.2 Psychometric Scaling Experiments

The user studies allowed us to assess that the proposed framework provides, with both the constrained and unconstrained strategies, an effective tool to remap parameters from a given source model to the parameters of an arbitrary target one. Further, they highlight that when there are noticeable differences between source and remapped renderings, this is due to deep differences between the characteristics of source and target models, *e.g.* remapping from Ward to GGX models (user studies 1 and 2). We further focus on this particular case and extend our validation with psychometric ranking experiments. The purpose of these is twofold:

- To provide validation that our framework provides the best achievable remapping, on average, when considering the parameter neighbourhood of the remapped values (*Perceptual Experiment 1*).
- To provide stronger evidence that the phenomena that causes the remapped renderings in the GGX model to be distinguishable from source renderings with parameters in the red portion of the plane $S_w \times R_w$ (see supplementary material), is actually the visually disruptive effect of the lack of Fresnel effect on one side (Ward), and its presence on the other side (GGX) (*Perceptual Experiment 2*).

We designed our experiments in Matlab, using the Psychophysics Toolbox Version 3 extensions [44], [45].

Perceptual Experiment 1. In the first experiment, subjects are presented with the source rendering on the right side of the screen and nine renderings on the left (Figure 15(a)). Only one rendering from 9 is generated using our remapped parameters, the others are rendered with parameters in the reflectance neighbourhood of the target parameter space. We focus on the plane given by $I_{ggx} \times R_{ggx}$, moving along one direction at a time.


The shapes are blobby objects, since they allow better discrimination of material reflectance than other shapes [46]. Their exact shape and order are randomised to avoid pixel-wise comparisons, rather than focusing on the overall reflectance. The diffuse value is fixed to a green shade, with values $RGB = [0.1, 0.3, 0.1]$. Observers were asked to select on the left side, the closest rendering to the source one in terms of surface properties ignoring the shape; once selected the image disappears and the observer must then choose the closest one from the eight left and so on, until all nine images are removed from the screen. For every

observer, this process was repeated 30 times for each of the selected 25 points in the Ward parameter space (on a 5×5 grid given by $S_w \times R_w$, $S_w = \{0.065, 0.13, 0.195, 0.26, 0.325\}$, $R_w = \{0.01, 0.05, 0.13, 0.29, 0.61\}$). The light probe used in this experiment was the ‘‘Courtyard of the Doge’s palace, Venice, Italy’’ [7]. Seven participants took part in this study, with each session lasting 2 to 3 hours.

Our rationale for this experiment is that on average among repetitions and observers, the visually closer target rendering to the source one, corresponding to the best set of target parameters for that given source, will have a lower ranking. This ranking should increase with the distance of the neighbours to the best one, giving rise to a ‘‘V’’ shape (centered on the latter) in a plot $neighbourID \times rank$. Flat areas indicate regions of the target parameter space, in which the renderings do not show perceivable differences with their immediate neighbours.

Figures 16(a) and 16(b) show the outcome of this experiment, with the red lines reporting the polynomial fit of measured data. For each of the plots on the x -axis we report the neighbouring identifiers, where ‘‘0’’ indicates the remapped parameters found by our framework (constrained strategy), of coordinates $\{I_0, R_0\}$ in the $I_{ggx} \times R_{ggx}$ space. The coordinate of the neighbours $\{-4, -3, \dots, 0, \dots, 4\}$ are reported in the supplementary material.

The analysis of the neighbours in the I_{ggx} direction shows that in 56% of the cases the observer preferred the rendering(s) corresponding to the parameters remapped by our framework (‘‘0’’), and in 37.5% of the cases its immediate neighbour ‘‘-1’’, indicating a preference towards a slightly lower value of the IOR and hence a slightly less visible Fresnel effect (Figure 16(e)). As for the roughness direction, our remapped parameters were preferred again in the 56% of the cases and ‘‘-1’’ in 28% of the cases. In many cases the ‘‘-1’’ neighbours have a very similar value of the fitness function (see Figure 16(d)), and the selection of $\{I_0, R_0\}$ by our framework is due to the constrained strategy.



(a) Perceptual Experiment 1

(b) Perceptual Experiment 2

Figure 15: Interface of the two psychometric scaling experiments.

Perceptual Experiment 2. In the second experiment, observers are presented with a simpler interface shown in Figure 15(b). At the bottom of the interface, a source BRDF rendering of a blobby shape is displayed. This image is used as a reference. At the top, two different renderings of blobby shapes are presented; one of them is rendered with the same model and parameters of the reference, while the second one is rendered using the remapped target parameters. In this forced-choice design the user must select, from the two shapes at the top, the most similar image to the reference in terms of surface properties.

The analysis of the results suggests that observers actually use the Fresnel effect as a cue to distinguish between source and remapped renderings, mainly located towards the silhouette of the objects, whereas other areas do not show significant differences;

this confirms the empirical finding from user studies 1 and 2. Please see supplementary materials for further details.

8 CONCLUSION

Frequently, digital artists face the problem of trying to match an image rendered somewhere else using a BRDF model that is unavailable within the rendering tool in use. Standard practice is to manually tweak a range of parameters for the available model, based on a subjective evaluation and experience, in a time-consuming fashion. In this study, we consider a real world scenario in which the source material model (BRDF) and parameters used to render a reference image are known. The BRDF available on the user side (target) is different from the source one. Despite having no access to the implementations of the material models, we propose an automatic solution to achieve consistent appearance across different models (and renderers).

We validate our framework through a series of experiments using both isotropic and anisotropic BRDF models. By employing 4 user studies and 2 psychometric scaling experiments, we are able to perform a detailed analysis of user perception of our remapping method. The results of these studies demonstrate that our parameter remapping framework allows us to obtain parameters, for the target model, which enable virtually indistinguishable renderings from those created with the source model, for a wide portion of the source model’s parameter space. This is true, even in challenging cases where the source and target models have extremely different characteristics, such as empirical models versus physically-based models with support for Fresnel reflectance. Further, our user studies also confirm the validity of our approach on complex scenes, which are more representative of a real world scenario.

Our experience with isotropic BRDF models provides a key insight; the differences between the statistical distribution of the microfacets used by the source and target models play a fundamental role in the quality achievable through remapping techniques.

REFERENCES

- [1] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, and T. Limperis, ‘‘Geometrical considerations and nomenclature for reflectance, nat,’’ *Bur. Stand. Rep., NBS MN-160*, 1977.
- [2] D. Guarnera, G. c. Guarnera, A. Ghosh, C. Denk, and M. Glencross, ‘‘Brdf representation and acquisition,’’ *Computer Graphics Forum*, vol. 35, no. 2, pp. 625–650, 2016.
- [3] R. Schregle, C. Denk, P. Slusallek, and M. Glencross, ‘‘Grand Challenges: Material Models in the Automotive Industry,’’ in *Eurographics Workshop on Material Appearance Modeling*, R. Klein and H. Rushmeier, Eds. The Eurographics Association, 2013.
- [4] A. Sztrajman, J. Krivanek, A. Wilkie, and T. Weyrich, ‘‘Image-based remapping of material appearance,’’ in *Proc. 5th Workshop on Material Appearance Modeling*, ser. MAM ’17, R. Klein and H. Rushmeier, Eds. Aire-la-Ville, Switzerland, Switzerland: The Eurographics Association, Jun. 2017.
- [5] A. Ngan, F. Durand, and W. Matusik, ‘‘Image-driven navigation of analytical brdf models,’’ in *Proceedings of the 17th Eurographics Conference on Rendering Techniques*, ser. EGSR ’06. Aire-la-Ville, Switzerland, Switzerland: Eurographics Association, 2006, pp. 399–407.
- [6] T.-W. Schmidt, F. Pellacini, D. Nowrouzezahrai, W. Jarosz, and C. Dachsbacher, ‘‘State of the art in artistic editing of appearance, lighting and material,’’ *Comput. Graph. Forum*, vol. 35, no. 1, pp. 216–233, Feb. 2016.
- [7] U.S.C. ICT. (2017) High-resolution light probe image gallery. [Online]. Available: <http://gl.ict.usc.edu/Data/HighResProbes/>
- [8] M. Colbert, S. Pattanaik, and J. Krivanek, ‘‘Brdf-shop: Creating physically correct bidirectional reflectance distribution functions,’’ *IEEE Computer Graphics and Applications*, vol. 26, no. 1, pp. 30–36, 2006.
- [9] X. An, X. Tong, J. D. Denning, and F. Pellacini, ‘‘Appwarp: Retargeting measured materials by appearance-space warping,’’ *ACM Trans. Graph.*, vol. 30, no. 6, pp. 147:1–147:10, Dec. 2011.

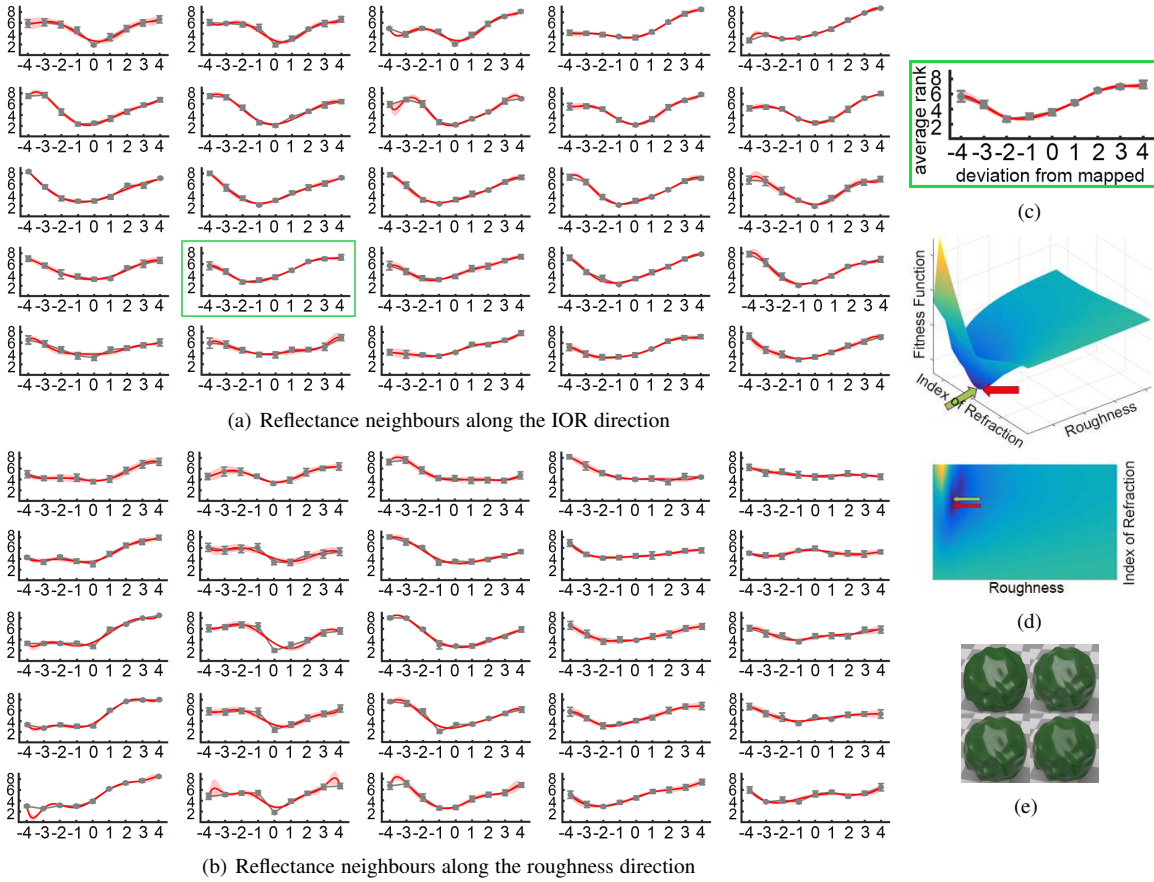


Figure 16: Psychometric ranking experiment. The 25 plots are arranged according to the grid $R_w \times S_w$, described in the text; the Ward roughness increases from left to right, the specularity from top to bottom. For each plot, the x -axis is the neighbour ID, and the y -axis the average ranking, as in (c), which shows a magnified version of the plot highlighted in green in (a). Some points in the reflectance neighbourhood have similar values for the fitness function (d). The green arrow points to the remapped values found by our framework, labelled “0” in (c), whereas the red arrows indicates the point “-2”. In (e), from left to right, top to bottom, renderings of the same object using the source model, neighbour “0” (our remapping), “-1” and “-2”. The renderings refer to the example in (c).

[10] A. Ngan, F. Durand, and W. Matusik, “Experimental analysis of brdf models,” in *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques*, ser. EGSR ’05. Aire-la-Ville, Switzerland, Switzerland: Eurographics Association, 2005, pp. 117–126.

[11] A. Brady, J. Lawrence, P. Peers, and W. Weimer, “genbrdf: Discovering new analytic brdfs with genetic programming,” *ACM Trans. Graph.*, vol. 33, no. 4, pp. 114:1–114:11, Jul. 2014.

[12] E. P. F. Laforge, S.-C. Foo, K. E. Torrance, and D. P. Greenberg, “Non-linear approximation of reflectance functions,” in *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH ’97. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1997, pp. 117–126.

[13] G. C. Guarnera, A. Ghosh, I. Hall, M. Glencross, and D. Guarnera, “Material capture and representation with applications in virtual reality,” in *ACM SIGGRAPH 2017 Courses*, ser. SIGGRAPH ’17. New York, NY, USA: ACM, 2017, pp. 6:1–6:72.

[14] M. Oren and S. K. Nayar, “Generalization of lambert’s reflectance model,” in *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*. ACM, 1994, pp. 239–246.

[15] M. M. Bagher, J. Snyder, and D. Nowrouzezahrai, “A non-parametric factor microfacet model for isotropic brdfs,” *ACM Trans. Graph.*, vol. 35, no. 5, pp. 159:1–159:16, Jul. 2016.

[16] A. Bousseau, E. Chapoulie, R. Ramamoorthi, and M. Agrawala, “Optimizing Environment Maps for Material Depiction,” *Proceedings of the Twenty-second Eurographics Conference on Rendering*, Eurographics Association, 2011, pp. 1171–1180.

[17] G. Wyszecki and W. S. Stiles, *Color science: Concepts and Methods, Quantitative Data and Formulae*. New York, USA: Wiley, 1982, vol. 8.

[18] M. Pedersen and J. Y. Hardeberg, “Full-reference image quality metrics: Classification and evaluation,” *Found. Trends. Comput. Graph. Vis.*, vol. 7, no. 1, pp. 1–80, Jan. 2012.

[19] A. Liu, W. Lin, and M. Narwaria, “Image quality assessment based on gradient similarity,” *IEEE Transactions on Image Processing*, vol. 21, no. 4, pp. 1500–1512, April 2012.

[20] B. T. Phong, “Illumination for computer generated pictures,” *Commun. ACM*, vol. 18, no. 6, pp. 311–317, Jun. 1975.

[21] G. J. Ward, “Measuring and modeling anisotropic reflection,” *SIGGRAPH Comput. Graph.*, vol. 26, no. 2, pp. 265–272, Jul. 1992.

[22] R. Rojas, *Neural Networks - A Systematic Introduction*. Springer Berlin Heidelberg, 1996.

[23] K. Deep, K. P. Singh, M. L. Kansal, and C. Mohan, “A real coded genetic algorithm for solving integer and mixed integer optimization problems,” *Appl. Mathematics and Computation*, vol. 212, no. 2, pp. 505–518, 2009.

[24] K. Deb, “An efficient constraint handling method for genetic algorithms,” *Computer methods in appl. mechanics and engineering*, vol. 186, no. 2-4, pp. 311–338, 2000.

[25] W. Jakob, “Mitsuba renderer,” 2010, <http://www.mitsuba-renderer.org>.

[26] G. J. Székely, M. L. Rizzo, and N. K. Bakirov, “Measuring and testing dependence by correlation of distances,” *The Annals of Statistics*, vol. 35, no. 6, pp. 2769–2794, 2007.

[27] M. Ashikhmin and P. Shirley, “An anisotropic phong brdf model,” *J. Graph. Tools*, vol. 5, no. 2, pp. 25–32, Feb. 2000.

[28] R. L. Cook and K. E. Torrance, “A reflectance model for computer graphics,” *ACM Trans. Graph. (TOG)*, vol. 1, no. 1, pp. 7–24, 1982.

[29] T. S. Trowbridge and K. P. Reitz, “Average irregularity representation of a rough surface for ray reflection,” *J. Opt. Soc. Am.*, vol. 65, no. 5, pp. 531–536, May 1975.

[30] B. Vogl. (2010). [Online]. Available: <http://dativ.at/lightprobes/>

[31] D. Geisler-Moroder and A. Dür, “A new ward brdf model with bounded albedo,” *Computer Graphics Forum*, vol. 29, no. 4, pp. 1391–1398, 2010.

[32] B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, “Microfacet models for refraction through rough surfaces,” in *Proceedings of the 18th Eurographics conference on Rendering Techniques*. Eurographics Association, 2007, pp. 195–206.

- [33] A. Sole, I. Farup, and P. Nussbaum, "Evaluating an image based multi-angle measurement setup using different reflection models," *Electronic Imaging*, vol. 2017, no. 8, 2017.
- [34] M. Toscani, M. Valsecchi, and K. R. Gegenfurtner, "Lightness perception for matte and glossy complex shapes," *Vision Research*, vol. 131, pp. 82–95, 2017.
- [35] J. Löw, J. Kronander, A. Ynnerman, and J. Unger, "Brdf models for accurate and efficient rendering of glossy surfaces," *ACM Trans. Graph.*, vol. 31, no. 1, pp. 9:1–9:14, Feb. 2012.
- [36] E. L. Church, P. Z. Takacs, "The prediction of brdfs from surface profile measurements," *Proc. SPIE* pp. 1165 – 1165 – 15, 1990.
- [37] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A data-driven reflectance model," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 759–769, Jul. 2003.
- [38] J. Dupuy, E. Heitz, J.-C. Iehl, P. Poulin, and V. Ostromoukhov, "Extracting microfacet-based brdf parameters from arbitrary materials with power iterations," *Computer Graphics Forum*, vol. 34, no. 4, pp. 21–30.
- [39] V. Havran, J. Filip, and K. Myszkowski, "Perceptually motivated brdf comparison using single image," *Computer Graphics Forum*, vol. 35, no. 4, pp. 1–12.
- [40] E. Heitz, "Understanding the masking-shadowing function in microfacet-based brdfs," *Journal of Computer Graphics Techniques*, vol. 3, no. 2, pp. 32–91, 2014.
- [41] R. K. Mantiuk, A. Tomaszewska, and R. Mantiuk, "Comparison of four subjective methods for image quality assessment," *Comput. Graph. Forum*, vol. 31, no. 8, pp. 2478–2491, Dec. 2012.
- [42] G. Turk and M. Levoy, "Zippered polygon meshes from range images," in *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '94. New York, NY, USA: ACM, 1994, pp. 311–318.
- [43] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '96. New York, NY, USA: ACM, 1996, pp. 303–312.
- [44] D. H. Brainard and S. Vision, "The psychophysics toolbox," *Spatial vision*, vol. 10, pp. 433–436, 1997.
- [45] M. Kleiner, D. Brainard, and D. Pelli, "What's new in psychtoolbox-3?" vol. 36, 2007, p. 1.
- [46] P. Vangorp, J. Laurijssen, and P. Dutré, "The influence of shape on the perception of material reflectance," *ACM Trans. Graph.*, vol. 26, no. 3, Jul. 2007.



Matteo Toscani Matteo Toscani studied Experimental Psychology in Florence, in the research group led by Professor David Burr. He moved to Giessen to work on his PhD at the Justus-Liebig-Universität. He gained his PhD on the relationship between lightness appearance and fixation behaviour in 2013. He is a postdoctoral fellow at Giessen University, working on material perception, fixation behaviour, and appearance in peripheral vision.



Mashhuda Glencross is director of research at Pismo Software. She has lectured at Leeds and Loughborough Universities, and worked as a product manager at ARM Ltd., and as a Research Fellow at The University of Manchester. Her work has spanned a diverse range of areas including; 3D reconstruction/imaging, material acquisition/modelling, visual perception, real-time/massive-model rendering, virtual reality and haptics.




Baihua Li is a senior lecturer in the Department of Computer Science and a member of Centre for Data Science at Loughborough University. She received her PhD degree in Computer Science from Aberystwyth University, and MSc and BSc degrees in Electronic Engineering from Tianjin University, China. Before joining Loughborough University, she worked at Manchester Metropolitan University.



Dar'ya Guarnera is a Research Associate at NTNU (Norway). She obtained her PhD in computer science from Loughborough University. Before this, she gained her Postgraduate Certificate in Education and first-class degree in Computer Science from Liverpool Hope University. She has a BSc in Architecture from Odessa State Academy of Civil Engineering and Architecture. Her interests include 3D architectural visualisation, virtual materials, art and sculpture.



Jon Yngve Hardeberg received his PhD from ENST, Paris. After a short industry career in the US, he returned to academia. He is Professor of Colour Imaging at the Norwegian Colour and Visual Computing Laboratory, NTNU. His research spans colour imaging, multispectral colour imaging, print and image quality, colorimetric device characterisation, material appearance, medical imaging, and cultural heritage imaging. He has co-authored around 300 research papers.



Giuseppe Claudio Guarnera is a Research Associate at NTNU (Norway). Prior to this, he worked as a Research Associate at Loughborough University. He received his PhD in computer vision and pattern recognition from the University of Catania (Italy). During his PhD, he also worked at the USC Institute for Creative Technologies (US). His research interests include colorimetry, virtual material acquisition, representation and human perception of materials.



Karl R. Gegenfurtner studied psychology in Regensburg (Germany). His PhD is in Experimental Psychology from New York University. After spending time as a PostDoc in New York and Tübingen, he became Professor of Psychology in Magdeburg. Since 2001, he has been at Giessen University (<http://www.allpsych.uni-giessen.de/karl/>).