

Perceptually Motivated BRDF Comparison using Single Image

V. Havran¹ J. Filip² K. Myszkowski³¹Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic²Institute of Information Theory and Automation of the CAS, Czech Republic³MPI Informatik, Saarbrücken, Germany

Abstract

Surface reflectance of real-world materials is now widely represented by the bidirectional reflectance distribution function (BRDF) and also by spatially varying representations such as SVBRDF and the bidirectional texture function (BTF). The raw surface reflectance measurements are typically compressed or fitted by analytical models, that always introduce a certain loss of accuracy. For its evaluation we need a distance function between a reference surface reflectance and its approximate version. Although some of the past techniques tried to reflect the perceptual sensitivity of human vision, they have neither optimized illumination and viewing conditions nor surface shape. In this paper, we suggest a new image-based methodology for comparing different anisotropic BRDFs. We use optimization techniques to generate a novel surface which has extensive coverage of incoming and outgoing light directions, while preserving its features and frequencies that are important for material appearance judgments. A single rendered image of such a surface along with simultaneously optimized lighting and viewing directions leads to the computation of a meaningful BRDF difference, by means of standard image difference predictors. A psychophysical experiments revealed that our surface provides richer information on material properties than the standard surfaces often used in computer graphics, e.g., sphere or blob.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1. Introduction

Modeling light reflectance from surfaces in a scene is an important component of realistic and predictive rendering. Bidirectional reflectance distribution functions (BRDF) capture the angular dependence for each incoming and each reflected light direction, which gives rise to complex 4D functions. The most convincing results are typically obtained using measured BRDFs, which might be captured with different angular resolutions and accuracy to accommodate the storage requirements and capturing costs. Moreover, such measured BRDFs are often further pre-filtered to reduce aliasing, lossy-compressed, or projected into rendering-efficient representations such as analytic functions [NDM05], precomputed polynomial [BAOR06, BAEDR08], tensor [SZC*07] and PCA [NJR15] bases. While all those steps are desirable, they might affect the accuracy of light reflection computation, which calls for a meaningful distance function between two BRDFs that accounts for such accuracy losses after the compression is applied.

In this work we propose a perceptually-motivated approach for the effective and meaningful difference computation between a pair of BRDFs. Our method just requires a single image to be rendered for the BRDF being tested; this is compared to the corresponding image for the reference BRDF using any of the existing quality

comparison methods. The key idea behind our method is the shape design of the rendered surface, which is optimized for the coverage in BRDF parameterization i.e. in terms of incoming and outgoing light directions. At the same time the shape of the surface needs to conform to our perceptual sensitivity to differences in material appearance (Fig. 1). The surface design is optimized altogether with the viewer and lighting positions to maintain the visibility of such BRDF samples by keeping shadowed and occluded surface parts under control. Our approach works well both for isotropic and anisotropic BRDFs, while a simple analytic surface formulation enables its efficient rendering. We conduct a series of psychophysical experiments that favorably compare the sensitivity of the human visual system (HVS) to the differences in the material appearance for our surface with respect to commonly used spheres and blobs [VLD07]. Furthermore, we demonstrate that the BRDF difference computation is far more reliable and discriminative for our surface when any of the state-of-the-art image quality metrics are used.

2. Related Work

In this section we discuss existing approaches for comparing BRDFs, which is the central issue addressed in this work. Since




Figure 1: Comparison of BRDFs using a single rendered image for different surfaces. We show that commonly used surfaces (spheres, blobs) struggle to achieve both: (1) a good visibility of differences (shown as psychometric measures in green) and (2) computational metric sensitivity (shown for ΔE in blue). The optimized surface on the right performs considerably better.

BTFs and SVBRDFs are not the main focus of this paper, we recall only those methods related spatially varying surface reflectance which encompass some perceptual components, or which have similar goals to ours in terms of optimizing for a wide coverage of incoming and outgoing light directions. Furthermore, we discuss surface design solutions which are driven by material appearance studies or by light reflection/refraction optimization. Finally, we briefly summarize perceptual findings which link material perception and shape.

BRDF distance computation – While analytical measures of BRDF similarity using Euclidean formulations are appealing due to their simplicity [PL07, FFGZ12], their predictions may not correlate well with actual perceived differences between rendered images for such BRDFs [NDM05]. For this reason image based metrics are typically employed, where often a simple object like a sphere is rendered under natural lighting [NDM05], as captured in high dynamic range (HDR) environment maps [Deb98]. While such natural lighting improves the discriminability between different materials [FDA03], it involves costly integration of incoming lighting over the hemisphere, and so in order to keep the computation efficient, a point light source [MAG*09] or a synthetically generated environment map [PR12] are often used. In the latter case the map characteristics are statistically consistent with natural lighting, but due to its radial symmetry and the isotropic BRDF assumption the integration can be efficiently performed. The final difference prediction in these approaches is computed using perceptual color metrics such as CIELAB [NDM05, PR12].

Recently, Filip [Fil15] has found that for anisotropic BRDFs a better discriminability of perceived differences between materials can be obtained when a point light source is used, which leads to less blurred light reflection patterns. In this work we also as-

sume an image-based approach, and since we consider fully 4D BRDF measurements we optimize for the reflective surface shape and a point light position to maximize the sample coverage in the incoming and outgoing light direction space. This naturally leads to the efficient computation of underlying images. With regard to the problem of isotropic BRDF discriminability, we attempt to improve it by imposing structure in the reflecting surface, which at least partially compensates for the desirable structure in natural illumination [FDA03].

BTF difference computation – Meseth et al. [MMK*06] evaluated the quality of images rendered using BTFs and compared them to corresponding photographs using objective image metrics and by performing a user study. Guthe et al. [GMSK09] developed an objective BTF difference metric, which employs a spatio-temporal contrast sensitivity function (ST-CSF) to prefilter BTF textures in spatial and angular domains prior to applying the standard CIE ΔE_{00} color difference equation. The proposed method proved useful in guiding the BTF compression. While in [MMK*06] the range of incoming and outgoing light directions is limited, Guthe et al. account for the complete BTF dataset under the assumption that the range of viewing conditions is plausible. In this work we have similar goals but with respect to BRDFs, which we achieve through a single image rendering under optimized lighting and viewing directions.

Filip et al. [FCGH08] run a psychophysical study and demonstrate that significant reduction of input BTF images is possible, while still providing the same visual quality. Jarabo et al. [JWD*14] run perceptual experiments to investigate the visual equivalence [RFWB07] of rendered images for different levels of BTF filtering, and find that blur in spatial domain is less tolerable than its angular counterpart. In this work, we investigate BRDF processing in angular domain, including filtering, but we emphasize on an objective BRDF difference computation.

Material design interfaces – In production software packages the so called material pickers are used for fine tuning the material appearance. Pellacini et al. [PFG00] advocate a perceptual linearization of physical parameters in reflectance models for their more intuitive usage, and through a multi-dimensional scaling (MDS) experiment they derive such a perceptual scaling for the Ward model parametrization. In the follow-up work, Kerr and Pellacini [KP10] show that in terms of efficiency in material tuning both physical parameters and their perceptually scaled counterpart are comparable. Ngan et al. [NDM06] propose an image-driven navigation over the space with embedded analytical BRDF models, in which the distance between the models is determined by one of the previously discussed BRDF difference methods [NDM05]. Wills et al. [WAKB09] extend this concept for measured BRDFs and propose a low-dimensional perceptual space for gloss, which enables an intuitive navigation and construction of new materials.

Spheres are commonly used in material pickers to provide visual feedback on the material appearance. Vangorp et al. [VLD07] show that discriminability thresholds for perceived differences between materials are lower for more complex shapes such as the bunny, buddha, or teapot than for the sphere, where the most consistent results across investigated materials and lighting have been

obtained for blobs. In the automobile industry a small-scale car-like reference shape (called the frog) is used to test paint finishes, as it features different curvatures similar to real cars. Ferwerda et al. [FWSP04] demonstrates that the frog is useful to render the car paint appearance as well. Nevertheless, neither the frog nor any of the shapes discussed here are optimized to systematically cover a wide range of incoming and outgoing light directions, which, as we show in this work, is crucial for image-based objective BRDF difference methods.

Reflective/refractive surface design – Precise guide of reflected illumination is required in many applications such as the design of automotive headlights, lamp fixtures, and backlight devices in displays [Elm80]. Optimization methods are typically applied in the modern design of reflective surfaces, which in case of free-form surfaces, e.g., defined by means of NURBS, might lead to complex optimizations due to a large number of control points that are required for precise guidance of such surface changes. To avoid this problem, suitable surface parameterizations that involve a limited number of parameters are often used [Fou10]. In our application such parametrization should provide sufficient control over the incoming and outgoing light directions coverage. The goals of our optimization are quite different than those in modern illumination applications, which typically concern maximization of light reflected in certain directions or its uniform distribution over the illuminated surface. We do not select faceted (tessellated) surface design [CDJ*10], which could give a more explicit control of the directions coverage, as it reduces the intuitiveness of surface shape for human observers and makes the material judgment more difficult [VLD07].

In graphics the design of surface shape to control the reflected and refracted lighting is an important research direction in material fabrication. Alexa and Matusik [AM10] fabricate a diffuse relief surface which represents two different images by changing the view direction, while Bermano et al. [BBAM12] employ relief self-shadowing to achieve similar effects. In goal-based caustics lens arrays are fabricated to reveal hidden image content but while looking through the lens a scrambled image version is seen [PHN*12] or desirable high contrast image projection can be achieved [STTP14]. While all those techniques involve the optimization of surface shape, the specific goals that guide such surface design are quite different than ours.

Material perception: Test surface design implications – Visual perception of materials is an active research field with extensive summaries presented in the recent survey paper [Fle14], where changes in the material perception due to interactions between lighting and object geometry are discussed. Olkkonen and Brainard [OB11] investigated gloss perception and found that an independent analysis of lighting and geometry does not lead to any solid prediction of their joint interactions. Nishida and Shinya [NS98] demonstrated that changes of the amplitude and spatial frequency in a bumpy surface limit the performance in the matching experiment where parameters in the Phong model are adjusted.

Nevertheless, surface curvature, which may vary greatly such as in the case of blobs, improves performance in material discrimination tasks [VLD07]. Such varying curvature reveals differ-

ent appearances of highlights, including compression in the highest curvature direction and stretching along the minimal curvature [NTO04]. The orientation structure of such specular patterns appears to be a powerful material characteristic information source in visual perception [FTA04]. For anisotropic BRDFs the direction of such stretches can be arbitrary and might depend on the material finishing structure, but perception of such materials has not been investigated so far. Since in this work we objectively compare anisotropic BRDFs our goal is essentially to best expose such patterns, and for this reason our test surface is composed of a number of differently oriented bumps with various curvatures, amplitudes, and scales. Also, bumpy surfaces receive more fixations than flat ones or regular gratings [FVH11], which might facilitate the perception of BRDF differences.

3. Overview

In this work we propose a BRDF distance function computing difference between two BRDFs. The proposed method fully relies on the comparison of a single rendered image against a reference image. Due to the single image requirement the design of the rendered scene should provide information on the material appearance which is as rich as possible both for the human observer as well as the objective image quality metrics. To this end, we apply optimization tools to derive the shape of the reflective surface, as well as the scene illumination and viewer position (Section 4). The surface covers all image pixels to maximize the number of independent BRDF samples, and the surface shape is optimized to achieve more uniform coverage of such samples in the BRDF parameter space, i.e. incoming and reflected light directions. We employ the optimized surface to investigate the HVS sensitivity for the detection of near-threshold BRDF differences as well as the magnitude estimation of larger BRDF differences. For this purpose we conduct a series of perceptual experiments, where additionally we consider other commonly used surfaces for the representation of material appearance such as spheres and blobs (Section 5). Based on the experiment outcome, we introduce further extensions in our surface design (Section 6), so that not only is the coverage of BRDF sampling optimized, but also the HVS sensitivity is improved; this is confirmed in another series of similar experiments (Section 7). Finally, we demonstrate a better performance of such optimized surfaces in BRDF difference detection and magnitude estimation using modern objective image quality metrics (Section 8).

4. Shape Optimization

Parametric surface – To optimize the coverage of incoming and reflected light directions one needs an analytic surface that allows its shape to be modified by using a limited number of parameters [Elm80]. Such a surface should account for convex, concave, and saddle curvatures of variable frequencies that are important for material appearance judgments [FWSP04] and provides meaningful input to the spatial vision in the HVS [MKRH11]. We tested a number of parametric surfaces and finally selected one shown in Eq. 1 defining surface 1. It is expressed in a polar space with axes $r \in (0, 1)$ and $\varphi \in (0, 2\pi)$.

$$S_1[z] = s[\sin(n_a \cdot \varphi)(r+1) \cos(n_r \cdot \pi \cdot r)] + \sqrt{R^2 - r^2} - R, \quad (1)$$

where s denotes the amplitude scale, R is the radius of the sphere that protrudes the test surface from below, n_r and n_a are the number of cycles along R and azimuth directions. In our optimization we assume the following ranges for the parameters s : [0.02–0.12], R : [1–12] meters, n_r : [2–9] and n_a : [3–15] cycles. The impact of parameter selection (within the specified ranges) on the surface appearance is demonstrated in Fig. 2. Note that impact of the parameter R is faint yet visible.





Figure 2: Effects of surface parameter selection in an orthographic projection: s – the amplitude scale, n_r – the number of cycles along radius, n_a – the number of cycles along the azimuth direction, and R – the radius of protruding sphere. The first three rows correspond to the minimal, average, and maximal parameter values in their respective ranges, while the last row shows extreme cases of parameter combinations.

Viewing and lighting conditions – The scene illumination should emphasize the differences in the material appearance so that any perceptible distortion in BRDF representation can be readily detected. This favors simple illumination with a single point light source, which best exposes high frequency features in reflected lighting including anisotropic patterns [Fil15]. Note that our task of BRDF difference detection is different than material recognition, in which case environment map lighting gives the most natural appearance [FDA03]. A single light point might lead to shadow regions, which do not provide any useful information on the BRDF characteristics and should be minimized.

Therefore, apart from four surface parameters, we also optimized the distances r_v, r_i and the directions $\theta_i, \phi_i, \theta_v, \phi_v$ to the light and camera, when we give their position in spherical coordinate system and the surface is located at the center of the coordinate system. We employ the perspective camera with the viewing angle of 50° and assume that the viewer distance to the center of an object falls into the range (0.3, 0.8) meters. Fig. 3 shows the effect of changes in the viewing distance r_v and light source distance r_i within the ranges that are used in the optimization.



(a) $r_v=0.3$ m (b) $r_v=0.8$ m (c) $r_i=1$ m (d) $r_i=12$ m

Figure 3: Effects of changes in the viewing distance: (a) $r_v = 0.3$ m, (b) $r_v = 0.8$ m, and the light source distance: (c) $r_i = 1$ m, (d) $r_i = 12.0$ m in a correct perspective projection.

Optimization procedure – Four surface parameters and six illumination/viewing parameters gives us a ten-dimensional optimization problem. We optimize these ten parameters to maximize the coverage of histogram of illumination and viewing angles, i.e., proportion of non-empty histogram bins. The histogram is a one-dimensional array that covers all illumination and viewing angles $\theta_i, \phi_i, \theta_v, \phi_v$, which are converted to the angles $\alpha_i, \beta_i, \alpha_v, \beta_v$ in the onion-sliced parameterization [HFM10] to avoid $(0, 2\pi)$ azimuthal discontinuity. These angles, all having the range of $(-\pi/2, \pi/2)$ were discretized to 13 bins in each dimension resulting in the histogram size of $(180/15 + 1)^4 = 13^4 = 28,561$ bins (each bin has the same surface area on a 4D hemisphere). In addition to maximizing coverage of histogram bins we also tested other statistical measures over the histogram. We have tried to optimize our parametric surface using the variance, total variation minimization of histogram bin values, and we tested also the Hellinger distance between the histogram bin values and a constant function. As the resulting surface shapes of these optimizations were similar for all those statistics, in this work we report the results for the histogram coverage in terms of incoming and reflected light directions, which we find quite intuitive (the ratio of histogram bins occupied).

The optimization was based on a simulated annealing algorithm implemented, for verification purposes, both in MATLAB and C++. The 4D histogram coverage was computed for rendering the surface shape for all lit pixels of an image $1,280 \times 960$ pixels. To suppress optimization results with excessive numbers of non-illuminated regions in the image, we set a threshold on the maximal percentage of such dark faces to 10% of all visible pixels. A typical running time of such optimization with 30,000 evaluations was 10 hours on an Intel Core i7-3610QM2.3GHz processor.

5. Experiment 1: Sensitivity to BRDF Distortions

Once the surface shape is optimized, its properties with respect to the discrimination of visible BRDF differences have to be compared with the baseline sphere and blob surfaces. To this end, we use BRDF filters that gradually degraded the original material appearance. We can then evaluate all tested surfaces by a direct comparison of original and degraded BRDFs.

5.1. BRDF Distortion Selection

We use three degradation filters, which represent typical inaccuracies in BRDF modeling and compression:

1. *smoothing* – filtering of azimuthal directions in BRDF using a box filter,




Figure 6: Materials used in our perceptual experiments: The first three are isotropic, while the remaining three are anisotropic.




Figure 4: Three distortion types as applied to the anisotropic fabric002 material. For each distortion type, the first row directly visualizes the BRDF data (the data presentation convention as in [FV14]) for five progressively increasing distortion levels, while the second row shows the distortion impact on the material appearance on a curved part of surface 1 (Fig. 5-right).

2. *anisotropic flattening* – a gradual transition between the original anisotropic appearance and its isotropic variant [Fil15],
3. *resampling* – down-sampling of the measured azimuthal directions.

Fig. 4 shows the effects of applying each of those filters on the entire BRDF and rendered surface. Note that for low distortion levels the effects are very faint, which might be challenging for existing difference image quality metrics. All of these filters were applied to all of the tested materials, except for anisotropic flattening which was only applied to the anisotropic materials.




Figure 5: Test surfaces used in Experiment 1. All scenes are the result of optimization, although for the sphere and blob only the viewing and light directions $\theta_i, \phi_i, \theta_v, \phi_v$ as well as the distances r_v, r_i were optimized.

5.2. Test Materials

We used six BRDFs from the UTIA BRDF Database (<http://btf.utia.cas.cz>) corresponding to different materials such as paint, leather, plastic, fabric, and wood. Three of them are isotropic (*carpaint01*, *leather03*, *plastic05*) and three are anisotropic (*fabric002*, *fabric005*, *wood42*) as shown in Fig. 6.

5.3. Psychophysical Study: Procedure

To evaluate the perceptual discriminability of BRDFs on the baseline and our optimized surface we conduct a psychophysical study. In this regard we prepared two types of experiments:

Threshold experiment – As experimental stimuli we have used pairs of static images of size 1920×1080 pixels, representing a material BRDF rendered on a surface. The experiment arrangement on the display was so that each pair consisted of the reference and degraded BRDFs as shown in Fig. 7-a (The images in Fig. 7 show the surface shape for experiment 2 described in the section 7). We used five levels of degradation for all distortion types (see Fig. 4). Pairs of images were displayed simultaneously, side-by-side in a random order. Three different surfaces were used; *sphere*, *blob*, and *surface1*. For each combination of material and distortion type there were 6 stimuli giving in total $6 \times (3 \text{ anisotropic} \times 2 \text{ filters} + 3 \text{ isotropic} \times 3 \text{ filters}) = 90$ stimuli (note that *anisotropic flattening* applies to anisotropic materials only). Subjects were asked to identify whether there is a visible difference between the two images.

The pairwise psychophysical data averaged across all subjects can be represented by the psychometric function $\psi(x)$ [WH01], which specifies the relationship between the underlying probability ψ of positive response and the stimulus intensity x : $\psi(x; \alpha, \beta, \gamma, \lambda) =$

$\gamma + (1 - \gamma - \lambda)F(x; \alpha, \beta)$, where F is the data fitting function with parameters α and β ; γ specifies the guess rate (i.e., response to the zero stimulus), and λ is the miss rate (i.e., a difference between incorrect and ideal responses for a large stimulus). We used Weibull function $F(x, \alpha, \beta) = 1 - \exp\left[-\left(\frac{x}{\alpha}\right)^\beta\right]$, where in our experiment configuration the correct response at the level of 50% probability determines the just-noticeable difference (JND) [VLD07].

Thurstonian scaling – As the threshold experiment does not provide us with a precise scaling of larger distortion magnitudes, in another experiment we employed Thurstonian scaling, i.e., measurement of the psychological scale separation between any two stimuli derived from Thurstone’s Law of comparative judgment [Thu27]. In our case, the scaling provides JND units of perceived difference [VMB*14]. We used a two-alternative forced choice (2AFC) experiment design with the reference BRDF shown on the display top (see Fig. 7-b). Subjects were asked to indicate which of the two bottom images was visually closer to the reference. As the bottom images can represent any combination of five degradation levels or the reference image, the final number of stimuli would be $5 \times 5 = 25$ per material and filter. This would expand to $25 \times (3 \times 2 + 3 \times 3) = 375$ stimuli. To make the experiment feasible we turned to an incomplete paired comparison design [SF01], and we restrict the comparison to two closest degradation levels, which reduces the number of stimuli to 255 images. To avoid a pixel-wise comparison, the upper reference image is always rendered for slightly rotated viewing angles (2°) [RFWB07].




Figure 7: Experiment 2: Stimuli examples (surface 3) for the threshold (a) and Thurstonian scaling (b) experiments.

Both experiments have been performed by 11 subjects; all had normal or corrected to normal vision, and were naive about the purposes of the experiment. Each subject used remote Nintendo Wii controllers to answer the questions, so that they could fully focus on the task. All stimuli were presented on a calibrated 24" LCD display (60 Hz, resolution 1920×1200 , peak luminance 120cd/m^2 , 6500K, gamma 2.2). The experiment has been performed under dim lighting conditions. Participants viewed the screen at a distance of 0.6 m.

5.4. Results

First, we analyzed the results of the psychophysical studies. The first row of Fig. 8 shows psychometric functions fitted to the results of pairwise comparisons averaged across all subjects and degradations. The error-bars represent the variance across all subjects. Here we can observe that subjects were on average more sensitive to the spherical surface while the optimized surface scored the least. Similar results were obtained also from the Thurstonian scaling in the second row of Fig. 8. Note that the derived JND scale is not absolute. By including a common point of reference (a reference image) we are able to link different scales for different distortions types, but strictly speaking we cannot compare scales between different surfaces. For such absolute comparisons psychometric functions can be considered as shown in the first row of Fig. 8.

We have also measured an average stimuli observation time, which generally decreases with the extent of BRDF degradation, i.e., the higher BRDF modifications were easier to spot. Also the comparison of psychometric functions for individual BRDF degradations revealed that effects introduced by the *sampling* filter were the most difficult to spot, while the subjects sensitivity to other filters was similar. The fastest response was achieved for the *fabric002* material that features significant anisotropic highlights, while the slowest response was measured for mostly diffuse *leather03*. Please refer to the supplemental materials for a more complete coverage of the response time measurements.

Fig. 10 shows the histogram coverage along with related statistics for the optimized scenes with the sphere, blob, and *surface 1*. The histograms are transformed into 2D images and show the bin coverage for the onion-sliced parameterization. Each small square block features histogram values dependent on β_v/α_v . The figure shows a significantly higher coverage of incoming and outgoing angles for our parameterized surface. Note that the sphere and blob were optimized only for illumination and viewing conditions, i.e., distances r_v, r_i and directions $\theta_i, \phi_i, \theta_v, \phi_v$ to the light and viewer. One can observe a three-fold increase of the coverage for the optimized surface when compared to the baseline surfaces.

6. Perception-motivated Shape Optimization

In the first experiment we learned that we can generate a surface with good histogram coverage but that if it is perceptually non-intuitive and relatively complex, it may distract observers and thus result in worse performance than baseline surfaces. Therefore, we decided to take the best of both worlds and combine our optimized surface function S_1 (Eq. 1) defining surface 1 with a hemisphere in the center, the size of which is comparable to the baseline sphere.




Figure 8: Experiment 1 for sphere, blob, and optimized surface 1. (a) psychometric functions, and (b) Thurstonian scaling.




Figure 9: Experiment 2 for surface 2 and surface 3. (a) psychometric functions, and (b) Thurstonian scaling.




Figure 10: Experiment 1: Comparison of histogram coverage for the sphere, blob, and the optimized surface 1.




Figure 11: Experiment 2: Comparison of histograms coverage of surfaces in the second experiment: optimized surfaces 2 and 3.

Details of the derivation of our new function S_2 defining surface 2 are given in the Appendix section. We use this combination of hemisphere and our analytic surface throughout the optimization procedure as described in Section 4, however we keep the hemisphere radius fixed. The outcome of such optimization is shown in Fig. 12-a-right.

Although this surface seems to be intuitively better than surface 1, the spiky bumps might be nonintuitive for some subjects. A surface which is even more familiar in appearance is to be preferred. To this end, we were considering a surface resembling a sphere with drapery lying over it. The folds of drapery can be approximated with the following function

$$f = s[10.0(r^2 \cdot \sin(n_a\phi + \sin(n_r\pi r))\sqrt{R^2 - r^2})], \quad (2)$$

where the meaning of parameters and the range of their values, as

used in the optimization, is the same as in Eq. 1, except that n_r denotes the number of zig-zag folds around the hemisphere. Fig. 12-b-left demonstrates the surface appearance as a function of the parameter selection that always remains within the ranges used in the optimization.

When we optimized this surface parameters for the coverage and combined it with a spherical surface we obtained the function S_3 defining surface 3 (see Appendix for details), which is visualized in Fig. 12-b-right. Final optimized parameters of all three custom surfaces are given in Appendix.

7. Experiment 2 – Combined Sphere and Optimized Surface

To understand the human performance in the discrimination of BRDF differences for the two optimized surfaces shown in Fig. 12-




Figure 12: Examples of (a) surface 2 and (b) surface 3 and their optimized versions on the right.




Figure 13: Experiment 2, eye tracking visualization: subjects gaze fixation hotspot-maps for both optimized surfaces.

right we repeated the threshold and Thurstonian scaling experiments as described in Section 5. A total of 13 subjects performed the experiments. Other settings and conditions remained the same. However, to learn what kind of visual information subjects use to perform their task, we have recorded eye movements for all of the subjects using the GazePoint GP3 eye-tracker. The device was calibrated for each subject individually and provided the locations and durations of fixations at a rate of 50 samples/s. The shortest fixation duration to be recorded was set to 10 ms.

7.1. Results

Fig. 9 shows the experiment outcome for the new surfaces. As can be seen the subject sensitivity is comparable to the baseline sphere surface (see Fig. 8), which was expected as the central hemisphere is now a part of the stimuli image.

The gaze fixation hotspots maps in Fig. 13, accumulated across 13 subjects, 6 materials, and 3 degradations, reveal that subjects used mainly the central spherical part for making their decision. The location of white peak correlates with the specular highlight and anisotropic singularity point at the pole of the hemisphere. Therefore, we assume that these visual features were the most important for subjects decisions. We can also observe that subjects did fixate surrounding geometry, especially for *surface 2*; this strategy was used for diffuse isotropic materials where the appearance of the hemisphere was not discriminative enough.

Fig. 11 shows that the histogram coverage for *surface 2* is signif-

icantly higher than *surface 1* (see Fig. 10) as well as *surface 3*. The improvement of coverage for *surface 2* is mainly due to its high frequency and gradient of peaks. While *surface 3* has the lower coverage, we assume that it is more intuitive for the observer.

Finally, we projected all incoming and outgoing directions present in the tested surfaces to half-angle parameterization [Rus98] and shown coverage histograms within characteristic slices of θ_h vs. θ_d that are often used for evaluation of isotropic BRDFs properties [NJR15]. Even, the projections in Fig. 14 clearly show that our optimized surfaces have higher coverage than the baseline surfaces.




Figure 14: Surfaces coverage histograms projected to elevation angles θ_h/θ_d representing characteristic BRDF slice of half-difference directional representation.

8. Objective Image Quality Metrics

The promising perceptual results of our optimized surfaces encouraged us to test the performance of various computational metrics on renderings of all of the tested surfaces. We considered ΔE , C-SSIM [LPU*13], and VDP2 [MKRH11] pixel-wise metrics (other metrics: RMSE, PSNR, SSIM are shown in the supplemental materials). Results in Fig. 15 show significantly higher sensitivity of the optimized surfaces compared to the baseline surfaces. We can observe similar sensitivity of the new surfaces (*surface 2*, *surface 3*) to the original *surface 1*. Behavior for individual materials is comparable for all optimized surfaces.

While ΔE and VDP2 had comparably high sensitivity to anisotropic materials (especially *fabric002* shown as red), C-SSIM has shown low sensitivity to material *fabric005*. This may be due to its dark color, which makes differences less visible. In contrast, while the isotropic materials *leather03*, *plastic05* exhibited very low sensitivity, the isotropic *carpaint01* material has shown sensitivity levels comparable to anisotropic ones. We expect that this is the result of its high dynamic range and the intensity of the specular peaks mean that these highlights are the most important visual feature regardless of the off-specular directions.

Our new surfaces are 2 to 3 times more sensitive than a sphere to pixel-wise computational metrics while still retaining a psychophysical response comparable to the baseline surfaces often used for comparison of material visualization fidelity in computer graphics. The subjects sensitivity for BRDF differences is comparable for both surfaces, while such difference discriminability by objective metrics is much better for *surface 2*. Please refer to the supplemental materials for for individual surface/material combinations.




Figure 15: Results of computational metrics for 6 BRDFs and all 5 tested surface shapes. ΔE (lower is more similar), C-SSIM (lower is more similar), and VDP2 (higher is more similar).

While our primary goal is not concerned with evaluating the performance of particular image quality metrics in predicting BRDF differences (which might depend on the metric itself as well as on the type of filter and magnitude of distortions, e.g., near threshold vs. supra-threshold distortions), we were interested in the correlation between subject responses and metric predictions as a function of the surface type. To this end, we investigated the Pearson and Spearman correlation between such responses as shown in Fig. 16. Overall, such correlation is slightly higher for our *surfaces 2 and 3* than for the sphere, while the blob and *surface 1* exhibit the lowest correlations. Remarkably, when the correlation is considered separately for anisotropic surfaces, it is higher than for isotropic surfaces. In particular, the correlation for *anisotropic flattening* is of the order of 0.9 and it is significantly higher than for other filter types, where it oscillates around 0.7 depending on the metric type. This might indicate that distortions of a clearly spatial nature, such as gradually filtered anisotropic patterns, are easier to detect by the existing quality metrics, which are often specifically trained for spatial distortions. The *smoothing* and *resampling* filtering is in fact performed in the angular domain of the BRDF, and while this translates into spatial patterns for smooth surfaces as considered in our experiments, the quality metrics are not explicitly designed to respond to such angular domain BRDF changes. Note that per pixel metrics such as RMSE and PSNR perform better in detecting such distortions (the *resampling* filter) than metrics with explicit spatial components such as VDP2 and SSIM. The specifics of the spatial component resulting from angular domain BRDF filtering do not involve significant contrast or structure changes that are captured by VDP2 and SSIM. This might indicate that specialized metrics need to be developed that somehow directly account for signal changes in the angular domain. Please refer to the supplemental materials for the complete set of correlation data.




Figure 16: Correlation between Thurstonian scaling data and objective metrics for all degradation filters and all materials: (a) Pearson and (b) Spearman correlation.

9. Discussions, Limitations and Extensions

In this paper we propose two optimized surfaces. One may wonder which one should be used in particular comparison applications. We recommend the use of *surface 2* for computational comparisons, due to its higher coverage that makes the pixel-wise comparison metrics more sensitive. On the other hand, the usage of *surface 3* is recommended, due to its better intuitiveness to visual applications, e.g., as BRDF selection/preview widget.

The main limitation of our method is that we cannot guarantee that the globally optimal surface has been found. One can possibly suggest other functions which could have even better coverage in the directional histogram. Moreover, for the sake of continuity and thus visual intuitiveness of the rendered surfaces, we restricted the variety of surfaces to those having less than 10% of pixels non-illuminated.

To support the usage of our results we provide code for visualizing our surfaces and for computing MSSIM differences of BRDFs

in C++. A list of pixel-wise angular illumination and viewing directions for each of the optimized surfaces is available for comparison and rendering of material appearance models (supplementary material). Unoptimized C++ code (single threaded, no SSE), as per the attachment to this publication, achieves 450 ms for computation of the BRDF distance using our method for image resolution 640×480 pixels. This could be further optimized both on a CPU and GPU as the SSIM computation can be highly parallelized. This approach considerably leverages the comparison process as it allows visualization of any material appearance model on our surfaces (for point light illumination) without the need to use any rendering tools.

The approach as presented is suitable for anisotropic BRDFs, however, it can easily be extended from a single image, as sufficient for anisotropic BRDFs, to a sequence of images for spatially varying BRDF and BTF data. By comparing a sequence of images we can draw a conclusion on the overall BTF dataset quality. The basic idea of this extension is that the BTF is moving on the proposed surface as piece of cloth over the shape. The average value of C-SSIM scores is computed for the sequence of image pairs to determine the difference between the reference and distorted BTFs (see the supplemental materials for more details). Fig. 17 shows an example of PCA compressed BTF data, when rendered for the sphere and Surface 2 and 3 scenes. The bottom row shows the corresponding difference maps.

Finally, although our surfaces were optimized for point-light illumination, they can be applied for more complex lighting scenarios such as environment maps (refer to the supplemental materials). In Fig. 18 we show an example of such complex lighting, and we demonstrate that for our surfaces the C-SSIM metric gains in the sensitivity to the fitting errors compared to the sphere when measured BRDFs are represented as analytic BRDF models [KSKK10].

10. Conclusions

This paper proposes a distance function between two BRDFs where the first BRDF is the reference and the second one is a distorted version of the reference which resulted from some kind of processing such as data compression. We show that the baseline surfaces usually employed for such comparisons (sphere, blob) are inefficient because their coverage of incoming and outgoing directions is very limited. Therefore we proposed the use of a number of alternative parametric surfaces and optimized their parameters using simulated annealing to achieve higher coverage of lighting directions. We show that such optimized surfaces achieve much higher sensitivity scores from pixel-wise computational metrics. Furthermore, we conducted psychophysical studies which proved that our surface allows the visibility of differences at least on a par with, if not better than, the standard baseline surfaces.

We assume that our optimized surface shapes will be helpful wherever a comparison of different models of surface reflectance is needed. We optimized our method on anisotropic BRDFs; however, our approach is easily extensible to comparing spatially-variable appearance (SVBRDF, BTF). We provide precomputed list of directions for each pixel of our surfaces, to allow fast material comparison without need of any rendering tools.

In future work we would like to investigate the use of arbitrary NURBS surfaces to achieve even higher flexibility of the optimized surface.

Acknowledgment

We would like to thank Rafał Mantiuk for providing us with the perceptual scaling software, Waqar Khan for help with the experiments, and Andrew Spence for the paper proofreading. Further we acknowledge the use of BTF data from BTF Database at the University of Bonn. This project was partially funded by the Czech Science Foundation under projects no. GA14-19213S and GA14-02652S, as well as the Fraunhofer and the Max Planck cooperation program within the framework of the German pact for research and innovation (PFI).

References

- [AM10] ALEXA M., MATUSIK W.: Reliefs as images. *ACM Trans. Graph. (Proc. SIGGRAPH)* 29, 4 (2010), 60:1–60:7. 3
- [BAEDR08] BEN-ARTZI A., EGAN K., DURAND F., RAMAMOORTHI R.: A precomputed polynomial representation for interactive BRDF editing with global illumination. *ACM Trans. Graph. (TOG)* 27, 2 (2008), 13:1–13:13. 1
- [BAOR06] BEN-ARTZI A., OVERBECK R., RAMAMOORTHI R.: Real-time BRDF editing in complex lighting. *ACM Trans. Graph. (Proc. SIGGRAPH)* 25, 3 (2006), 945–954. 1
- [BBAM12] BERMANO A., BARAN I., ALEXA M., MATUSIK W.: Shadowpix: Multiple images from self shadowing. *Comp. Graph. Forum* 31, 2p3 (2012), 593–602. 3
- [CDJ*10] CASSARLY W. J., DAVID S. R., JENKINS D. G., RISER A. P., DAVENPORT T. L.: Automated design of a uniform distribution using faceted reflectors. *Optical Engineering* 39, 7 (2010), 1830–1839. 3
- [Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. ACM SIGGRAPH* (1998), pp. 189–198. 2, 11
- [Elm80] ELMER W. B.: *The Optical Design of Reflectors*. Wiley, 2nd edition, 1980. 3
- [FCGH08] FILIP J., CHANTLER M. J., GREEN P. R., HAINDL M.: A psychophysically validated metric for bidirectional texture data reduction. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 27, 5 (2008), 138:1–138:11. 2
- [FDA03] FLEMING R. W., DROR R. O., ADELSON E. H.: Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3, 5 (2003), 3. 2, 4
- [FFGZ12] FORES A., FERWERDA J., GU J., ZHAO X.: Toward a perceptually based metric for BRDF modeling. In *20th Color and Imaging Conference* (2012), CIC'12, pp. 142–148. 2
- [Fil15] FILIP J.: Analyzing and predicting anisotropic effects of BRDFs. In *ACM SAP* (2015), pp. 25–32. 2, 4, 5
- [Fle14] FLEMING R. W.: Visual perception of materials and their properties. *Vision Research* 94, 0 (2014), 62 – 75. 3
- [Fou10] FOURNIER F.: *Freeform Reflector Design with Extended Sources*. PhD thesis, University of Central Florida, 2010. 3
- [FTA04] FLEMING R. W., TORRALBA A., ADELSON E. H.: Specular reflections and the perception of shape. *Journal of Vision* 4, 9 (2004), 10. 3
- [FV14] FILIP J., VÁVRA R.: Template-based sampling of anisotropic BRDFs. *Comp. Graph. Forum (Pacific Graph)* 33, 7 (2014), 91–99. 5




Figure 17: Examples of BTF rendered on the sphere and optimized surfaces 2 and 3, together with C-SSIM luminance-contrast maps depicting differences due to using 20 PCA components instead of 60 (material knitted wool). Material wool, sphere: CSSIM = 0.0134, surface 2: CSSIM=0.0531, surface 3: CSSIM= 0.0286. The Wool material BTF courtesy of the University of Bonn.




Figure 18: Comparison of the reference BRDFs for the carpaint01 and fabric005 materials (the top row) and their analytic approximations [KSKK10] (the bottom row) for different surfaces as illuminated by the grace environment map [Deb98]. The numerical values under the bottom row are the difference measures obtained with the C-SSIM metric (lower is more similar).

[FVH11] FILIP J., VACHA P., HAINDL M.: Analysis of human gaze interactions with texture and shape. In *S+SSSPR Workshop, Springer LNCS 7252* (2011), pp. 160–172. 3

[FWSP04] FERWERDA J. A., WESTIN S. H., SMITH R. C., PAWLICKI R.: Effects of rendering on shape perception in automobile design. In *APGV* (2004), pp. 107–114. 3

[GMSK09] GUTHE M., MÜLLER G., SCHNEIDER M., KLEIN R.: BTF-CIElab: A perceptual difference measure for quality assessment and compression of BTFs. *Comp. Graph. Forum* 28, 1 (2009), 101–113. 2

[HFM10] HAVRAN V., FILIP J., MYSZKOWSKI K.: Bidirectional texture

function compression based on the multilevel vector quantization. *Comp. Graph. Forum* 29, 1 (2010), 175–190. 4

[JWD*14] JARABO A., WU H., DORSEY J., RUSHMEIER H., GUTIERREZ D.: Effects of approximate filtering on the appearance of bidirectional texture functions. *IEEE TVCG* 20, 6 (2014), 880–892. 2

[KP10] KERR W. B., PELLACINI F.: Toward evaluating material design interface paradigms for novice users. *ACM Trans. Graph. (Proc. SIGGRAPH)* 29, 4 (2010), 35:1–35:10. 2

[KSKK10] KURT M., SZIRMAY-KALOS L., KRÍVÁNEK J.: An anisotropic BRDF model for fitting and monte carlo rendering. *SIGGRAPH Comput. Graph.* 44 (2010), 3:1–3:15. 10, 11

- [LPU*13] LISSNER I., PREISS J., URBAN P., LICHTENAUER M. S., ZOLLIKER P.: Image-difference prediction: From grayscale to color. *IEEE TIP* 22, 2 (2013), 435–446. 8
- [MAG*09] MATUSIK W., AJDIN B., GU J., LAWRENCE J., LENSCH H. P. A., PELLACINI F., RUSINKIEWICZ S.: Printing spatially-varying reflectance. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 28, 5 (2009), 128:1–128:9. 2
- [MKRH11] MANTIUK R., KIM K. J., REMPEL A. G., HEIDRICH W.: HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Trans. Graph. (Proc. SIGGRAPH)* 30, 4 (2011), 40:1–40:14. 3, 8
- [MMK*06] MESETH J., MÜLLER G., KLEIN R., RÖDER F., ARNOLD M.: Verification of rendering quality from measured BTFs. In *APGV* (2006), pp. 127–134. 2
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. In *EGSR* (2005), pp. 117–126. 1, 2
- [NDM06] NGAN A., DURAND F., MATUSIK W.: Image-driven navigation of analytical BRDF models. In *EGSR* (2006), pp. 399–407. 2
- [NJR15] NIELSEN J. B., JENSEN H. W., RAMAMOORTHY R.: On optimal, minimal BRDF sampling for reflectance acquisition. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 34, 6 (2015), 186:1–186:11. 1, 8
- [NS98] NISHIDA S., SHINYA M.: Use of image-based information in judgments of surface-reflectance properties. *J. Opt. Soc. Am. A* 15, 12 (1998), 2951–2965. 3
- [NTO04] NORMAN J. F., TODD J. T., ORBAN G. A.: Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychological Science* 15, 8 (2004), 565–570. 3
- [OB11] OLKKONEN M., BRAINARD D. H.: Joint effects of illumination geometry and object shape in the perception of surface reflectance. *i-Perception* 2, 9 (2011), 1014. 3
- [PFG00] PELLACINI F., FERWERDA J. A., GREENBERG D. P.: Toward a psychophysically-based light reflection model for image synthesis. In *Proc. ACM SIGGRAPH* (2000), pp. 55–64. 2
- [PHN*12] PAPAS M., HOUIT T., NOWROUZEZAHRAI D., GROSS M., JAROSZ W.: The magic lens: Refractive steganography. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 31, 6 (2012), 186:1–186:10. 3
- [PL07] PELLACINI F., LAWRENCE J.: Appwand: Editing measured materials using appearance-driven optimization. *ACM Trans. SIGGRAPH* 26, 3 (2007). 2
- [PR12] PEREIRA T., RUSINKIEWICZ S.: Gamut mapping spatially varying reflectance with an improved BRDF similarity metric. *Comp. Graph. Forum* 31, 4 (2012). 2
- [RFWB07] RAMANARAYANAN G., FERWERDA J., WALTER B., BALAK.: Visual equivalence: Towards a new standard for image fidelity. *ACM Trans. Graph. (Proc. SIGGRAPH)* 26, 3 (2007). 2, 6
- [Rus98] RUSINKIEWICZ S.: A new change of variables for efficient BRDF representation. In *EGSR* (1998), pp. 11–22. 8
- [SF01] SILVERSTEIN D. A., FARRELL J. E.: Efficient method for paired-comparison. *J. El. Imaging* 10, 2 (2001), 394–398. 6
- [STTP14] SCHWARTZBURG Y., TESTUZ R., TAGLIASACCHI A., PAULY M.: High-contrast computational caustic design. *ACM Trans. Graph. (Proc. SIGGRAPH)* 33, 4 (2014), 74:1–74:11. 3
- [SZC*07] SUN X., ZHOU K., CHEN Y., LIN S., SHI J., GUO B.: Interactive relighting with dynamic BRDFs. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 26, 3 (2007). 1
- [Thu27] THURSTONE L. L.: A law of comparative judgment. *Psychological Review* 34 (1927), 273–286. 6
- [VLD07] VANGORP P., LAURIJSEN J., DUTRÉ P.: The influence of shape on the perception of material reflectance. *ACM Trans. Graph. (Proc. SIGGRAPH)* 26, 3 (2007), 77. 1, 2, 3, 6

- [VMB*14] VANGORP P., MANTIUK R. K., BAZYLUK B., MYSZKOWSKI K., MANTIUK R., WATT S. J., SEIDEL H.-P.: Depth from HDR: Depth induction or increased realism? In *ACM SAP* (2014), pp. 71–78. 6
- [WAKB09] WILLS J., AGARWAL S., KRIEGMAN D., BELONGIE S.: Toward a perceptual space for gloss. *ACM Trans. Graph. (TOG)* 28, 4 (2009), 103:1–103:15. 2
- [WH01] WICHMANN F., HILL N.: The psychometric function: I. fitting, sampling, and goodness of fit. *Perception & Psychophysics* 63, 8 (2001), 1293–1313. 5

Appendix

Analytical functions combining hemisphere and the optimized surfaces as used in the second experiment.

Surface 2:

$$S_2[z] = s[\text{TX}(r, t_1, t_2) \cdot \sin(n_a 2\phi) \cdot \cos(m_1(r, r_1, t_1, n_r)) + \text{HS}(r, 0.4)(1.0 - \text{TX}(r, t_1, t_2))] + \sqrt{R^2 - r^2} - R, \quad (3)$$

where $t_1 = 1.0k$, $t_2 = 1.5k$, and $k = 0.09$ is the radius of hemisphere included by function $\text{HS}(\cdot)$ shown below. Functions $\text{TX}(\cdot)$ and $m_1(\cdot)$ take care for smooth transition between hemisphere and the optimized surface. Surface parameters are s amplitude scale, n_r cycles per radius, n_a cycles per azimuth, and R is radius of the sphere protruding the test surface from below.

Surface 3:

$$S_3[z] = s[10m_2(r, s)^2 \sin(n_a \phi + \sin(n_r \pi m_2(r, s))) + \sqrt{R^2 - m_2(r, s)^2}] + \text{HS}(r, 0.25), \quad (4)$$

where parameters are the same as above except n_r representing number of zig-zag moves of the folds across surface diameter (2-9) and the function $m_2(\cdot)$ just clamps the radius.

```

function  $r_u = \text{HS}(r, c)$ 
1:  $k = 0.09$  // sphere radius - fixed size
2: if ( $r < k$ )  $r_u = \sqrt{k^2 - r^2}$ 
3: else if ( $r < 2k$ )  $r_u = -\sqrt{k^2 - (2k - r)^2}$ 
4: else  $r_u = -k$ 
function  $v = \text{TX}(x, t_1, t_2)$ 
1:  $v = (x - t_1) / (t_2 - t_1)$  // smooth transition between
2: if ( $v < 0$ )  $v = 0$  //  $t_1$  and  $t_2$ 
3: if ( $v > 1$ )  $v = 1$ 
4:  $v = v^3(v(6v - 15) + 10)$  // output is 0 to 1
function  $r = m_1(r, r_1, t_1, n_r)$ 
1: if ( $r \geq t_1$ )  $r = ((r - t_1)n_r + 0.25)\pi$ 
2: else  $r = 0$ 
function  $r = m_2(r, s)$ 
1: if ( $r \leq s$ )  $r = 0$ 

```

Parameters of the optimized surfaces S

S	s	n_r	n_a	R	r_v	r_i	θ_v	ϕ_v	θ_i	ϕ_i
	[-]	[-]	[-]	[m]	[m]	[m]	[°]	[°]	[°]	[°]
1	0.060	8.0	5.0	4.3	0.8	6.7	29.4	148.5	28.5	277.1
2	0.091	14.5	6.0	5.2	0.7	4.6	46.2	249.2	3.7	51.7
3	0.047	15.0	9.0	1.9	0.7	2.5	36.3	209.7	19.4	14.3