DOI: 10.1111/cgf.12789

COMPUTER GRAPHICS *forum*
Volume 00 (2016), number 00 pp. 1–12

Predicting Visual Perception of Material Structure in Virtual Environments

J. Filip, R. Vávra, M. Havlíček and M. Krupička

Institute of Information Theory and Automation of the Czech Academy of Sciences, Czech Republic
{jiri.filip, vavra, havlimi2, krupimik}@utia.cas.cz

Abstract

One of the most accurate yet still practical representation of material appearance is the Bidirectional Texture Function (BTF). The BTF can be viewed as an extension of Bidirectional Reflectance Distribution Function (BRDF) for additional spatial information that includes local visual effects such as shadowing, interreflection, subsurface-scattering, etc. However, the shift from BRDF to BTF represents not only a huge leap in respect to the realism of material reproduction, but also related high memory and computational costs stemming from the storage and processing of massive BTF data. In this work, we argue that each opaque material, regardless of its surface structure, can be safely substituted by a BRDF without the introduction of a significant perceptual error when viewed from an appropriate distance. Therefore, we ran a set of psychophysical studies over 25 materials to determine so-called critical viewing distances, i.e. the minimal distances at which the material spatial structure (texture) cannot be visually discerned. Our analysis determined such typical distances typical for several material categories often used in interior design applications. Furthermore, we propose a combination of computational features that can predict such distances without the need for a psychophysical study. We show that our work can significantly reduce rendering costs in applications that process complex virtual scenes.

Keywords: appearance modelling, reflectance & shading models, user studies, BRDF, BTF

ACM CCS: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Colour, shading, shadowing and texture

1. Introduction

Nowadays, photorealistic renderings of real-world materials are widely used in computer graphics applications spanning from the motion picture industry, computer games or paint industry to visual safety simulations or virtual prototyping in the automotive and aircraft industry, or architecture. Such renderings rely on various representations of material appearance that can offer various levels of realism, but unfortunately also different data acquisition, storage and related computational costs. In most static virtual scenes, the data storage and computational costs can be reduced by means of selecting a minimal material appearance representation that can still introduce material properties sufficiently at minimal resource costs. This paper focuses specifically on the analysis and modelling of the human ability to distinguish material structure (or texture) in different material categories. We consider this information important to estimate a distance at which the material structure becomes visually indistinguishable. Figure 1 illustrates this effect on a set of real photographs of material observed from distances increasing from 0.6 to 4.8 m.

As we target our work on virtual environments, we analyse human visual perception of material structure directly on digitized representations of material appearance.

Probably the most common representation of material appearance is the *Bidirectional reflectance distribution function* (BRDF) [NRH*77] $BRDF(\theta_i, \varphi_i, \theta_v, \varphi_v)$. The BRDF represents reflectance of opaque materials dependent on local illumination $\mathbf{I}(\theta_i, \varphi_i)$ and view $\mathbf{V}(\theta_v, \varphi_v)$ directions often represented by spherical angles.

While a four-dimensional BRDF describes the distribution of energy reflected to the viewing direction when illuminated from a specific direction, a six-dimensional *bidirectional texture function* (BTF) [DvGNK99] $BTF(x, y, \theta_i, \varphi_i, \theta_v, \varphi_v)$ additionally captures the spatial dependency $[x, y]$ of reflectance across a material surface.

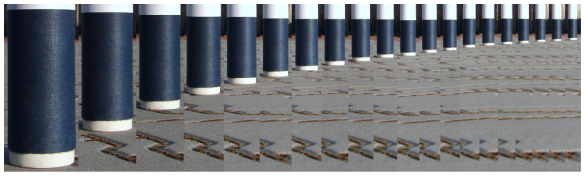

Figure 1: *A photo-collage of real non-woven fabric material observed from distances ranging between 0.6 and 4.8 m (step 0.2 m).*

While the BRDF using moderate directional resolution in HDR quality occupies, e.g. $3 \times 100 \times 100 = 30\,000$ float values, the BTF of the same angular and very low spatial resolution 128^2 takes about $128 \times 128 \times 3 \times 100 \times 100 = 491\,520\,000$ values. This tremendous difference in storage size implies higher memory and computational demands in related rendering applications.

While BRDF imposes restrictions on reciprocity and opacity, the six-dimensional BTF generally does not fulfil these restrictions. This is due to local effects in a rough material structure such as occlusions, masking, subsurface scattering and interreflections. These effects constitute a great difference in BTF realism as compared to BRDF. However, if we move away from the viewed surface, the spatial structure and local effects become less apparent. Finally at some viewing distance we arrive at an appearance that is equivalent to a BRDF, i.e. exhibiting the appearance of a homogeneous, flat surface. This is common in remote sensing where, e.g. foliage or urban areas appearances are, due to long viewing distance, represented by a BRDF [QKM*00, SGS*02]. Therefore, we argue that all opaque materials can be represented by a BRDF when viewed from an appropriate distance without compromising its basic properties.

Thus the main contributions of this paper are:

- A psychophysical analysis of structure visibility for common material categories, identifying distances when the structure becomes indistinguishable for the human observer.
- Analysis of factors that define material structure visibility allowing its statistical prediction for arbitrary unknown materials.
- Analysis of rendering speed gain obtained when our model of structure visibility was applied for on-the-fly selection of material appearance representation.

The paper is structured as follows. Section 2 sets our work into a context of related works in the field. Section 3 describes materials tested and provides technical information of the psychophysical studies. Section 4 summarizes results of the studies and Section 5 use them to estimate computational features predicting the results. Section 6 shows rendering application of the results. Section 7 concludes the paper and discusses future challenges.

2. Related Work

Our work is related to the realistic representation of materials in a virtual environment; however, partly deals with the reliable prediction of material structure visual perception. We discuss the related work separately.

Psychophysical analysis of materials in virtual environments– Although visual perception of material appearance represented by means of isotropic BRDFs has been thoroughly studied as a function of global illumination and 3D geometry [VLD07, RFWB07, KFB10] the perceptual effects of BTFs were not studied in such a width to date. Although the material appearance by means of BTF has been a subject of much research (see [FH09] for a review), most of the analysis has been focused on the compression and rendering of these massive data sets. The psychophysical analysis of BTFs has been limited either to the verification of visual quality [MMK*06] or for the guiding of data compression algorithms [FCGH08, FHC10]. Guthe *et al.* [GMSK09] applied model of achromatic contrast sensitivity function (CSF)[BK80] to achieve a more effective BTF compression without significant loss of perceived fidelity. The recent work of Jarabo *et al.* [JWD*14] psychophysically analyses effect of filtering (downsampling and oversampling) in both spatial and directional domains. Although this paper is very relevant to our work, its conclusions are more general and the way the studies were designed does not help us to derive the required material-dependent critical distance.

Structure visibility prediction– Our work also closely relates to applied aspects of human visual system (HVS) and its models. The contrast sensitivity of HVS was extensively studied and modelled in the past [Wan95]. The spatiotemporal visibility of texture was modelled in [BK80] by a product of a spatial and temporal frequency response curves. These contrast sensitivity curves are approximated by the model based on two space–time separable Gaussian filters. Furthermore, luminance and chromatic contrast of stimuli images can be estimated directly from pixel-wise cone responses to stimuli images (cone channels, respectively) according to [WM97]. Alternatively, image salience [PN04] can be used. This method predicts visual fixations by combination of first- and second-order image statistics, namely, luminance and texture contrasts based on per-pixel spatial gradients over stimuli luminance represented by a Gaussian-Laplacian pyramid. Another more applied approaches predicts visual difference between two images either by a simple statistical model [WBSS04] or by means of a more involved modelling of low-level visual perception [MKRH11].

Although this research provided interesting insights into human perception of digitized materials appearance, there are still unresolved connections between high-level human perception of general material appearance and its reliable computational models. One of such areas, that can potentially simplify materials rendering methods, is an assessment of the material-dependent perceptual effects of viewing distance. Therefore, we designed a set of psychophysical studies to detect the so called *critical viewing distance*, i.e. viewing distance at which the structure of the material becomes impossible to identify.

3. The Experiments

This section describes materials used in our analysis, their digital representation used, test scenes and details of the three performed psychophysical studies.

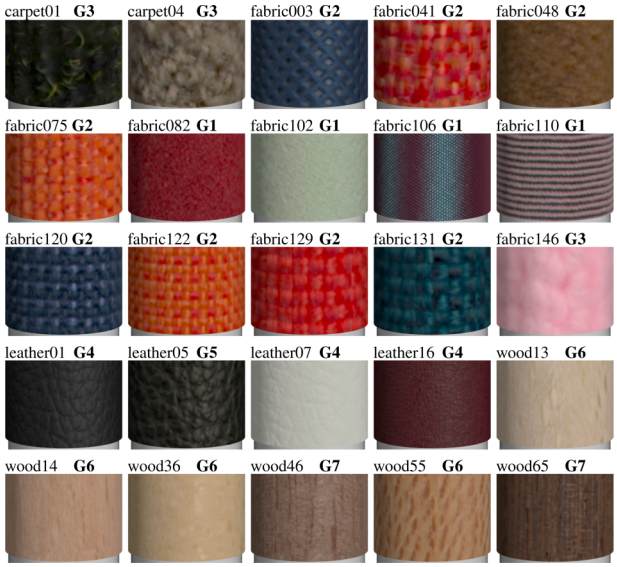

Figure 2: *Overview of 25 materials used in our study (listed alphabetically) and their distribution into seven analysed groups.*

3.1. Tested materials

We tested 25 materials out of categories often used in interior design (carpet, fabric, leather and wood). Their BTFs were measured at an angular resolution of 81 × 81 for both illumination and viewing directions [SSK03], and a spatial resolution of 350 DPI. Examples of materials are shown in Figure 2. For the sake of seamless coverage of the test objects in required spatial resolution, we estimated the single seamless repeatable tile using approach [SH05]. Sizes of tiles ranged from 80 to 300 pixels. As some of the tiles were not entirely spatially uniform in hue and luminance, we carefully removed their lowest frequency components in Fourier space. This step avoids visually distractive tile repetition, that could introduce low frequencies that are not present in the original material. The approximate BRDFs of materials were obtained by spatial averaging across tile area. We consider this approximation as reasonably accurate as size of the biggest tiles was 22 mm, which is negligible when compared to the distance between the sample and light/camera which is 1.3/2.1 m.

Next, we visually estimated the physical size of the biggest structure elements of the tested materials. The results in millimetres are shown in Figure 3 and overlaid over materials in the Supporting Information. We further analysed height of the materials by integration of normals resulting from overdetermined photometric stereo using 80 illumination directions (frontal illumination was removed to reduce presence of specular reflections). Although one could argue that we could compute local facets slope distributions variance for determining the size of material structure, we consider this approach as unreliable especially for translucent mesostructure fabric materials. For the sake of further analysis, we divided materials into seven categories based on the type of material, its structure size and height:

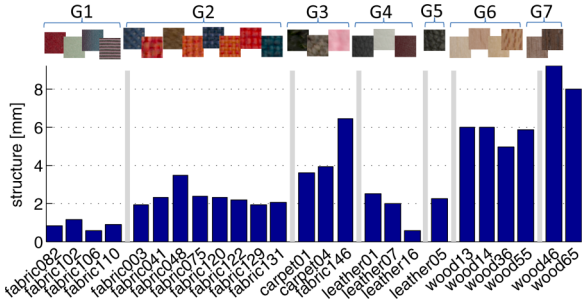

Figure 3: *Categorization of the studied materials into seven groups based on a size of their biggest structure element.*

G1: *fabric smooth* group comprises materials *fabric082*, *fabric102*, *fabric106* and *fabric110*. This group represents mostly apparel fabric materials with a fine structure of typical size of the biggest structure element being below 1 mm.

G2: *fabric meso* group comprises materials *fabric003*, *fabric041*, *fabric048*, *fabric075*, *fabric120*, *fabric122*, *fabric129* and *fabric131*. This group represents typical upholstery materials used (e.g. on office chairs) with the size of structure element being below 3 mm.

G3: *fabric rough* comprises materials *carpet01*, *carpet04* and *fabric146*. It represents very rough fabric material: office carpet and a crocheted sweater with the structure size above 4 mm.

G4: *leather meso* group comprises materials *leather01*, *leather07* (both synthetic) and *leather16* (genuine). This category comprises relatively flat leather materials (often used in car interior design or upholstery) with the structure size of up to 2 mm.

G5: *leather rough* comprises single rough genuine leather material *leather05*. We designed a special category for this material due to its higher surface height that sets its appearance apart from leather materials in the previous group.

G6: *wood meso* group includes *wood13* (beech), *wood14* (steamed beech), *wood36* (european lime) and *wood55* (plane). These wood materials exhibit smooth surfaces, smaller structure elements (below 6 mm) and less contrasted appearance.

G7: *wood rough* group include material *wood46* (american walnut) and *wood65* (wenge). This group represents contrasting wood samples with characteristic and long vertical grooves of sizes above 8 mm.

We also included into the experiment stimuli two artificial chessboard patterns of square width being 2 and 5 mm. While the first one was used to check subjects' resolution, the second one was used for filtering subjects who either have impaired vision or did not understand the task.

3.2. Test scene

As the basic notion of our study is the psychophysical detection of a critical viewing distance where the structure becomes visually indistinguishable, we altered the task into an analysis of the difference between a BRDF and a BTF as a function of viewing distance. Therefore, we generated a virtual scene consisting of a

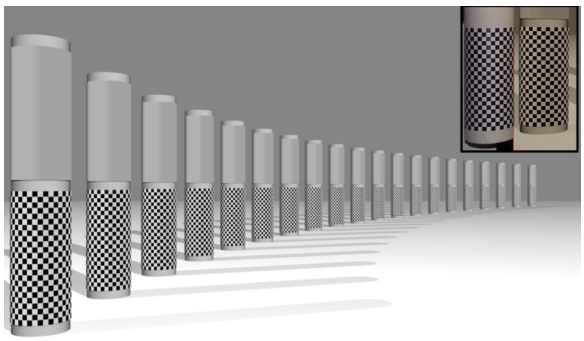

Figure 4: *Example of the virtual scene consisting of set of consequently receding cylinders showing the BRDF and BTF of the same material. Top-right: a side-by-side comparison of real cylinder (left) with its representation on the screen (right) both having the same width and 5 mm checkerboard pattern.*

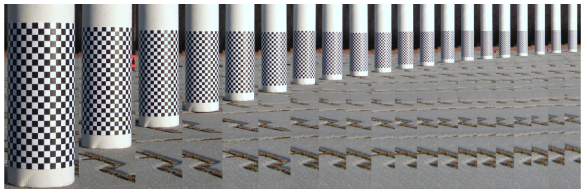

Figure 5: *The virtual scene validation using a set of photos of the physical cylinder from the same distances as in Figure 4.*

set of consequently receding cylinders as shown in Figure 4. The upper part of the cylinders shows a BRDF while the lower part shows a BTF of the same material. We opted for illumination by means of two directional light sources; one from the right (intensity 1.0) and the second from the left (intensity 0.3). Such configuration guarantees the visibility of structure on all visible areas of cylinders and preserves contrast shadowing in the material structure that would be smoothed by a typical environment illumination. We selected a cylinder due to its clearly defined texture mapping without the need of warping or cutting of the mapped material. The perspective projection of the scene is calibrated so as its appearance resembles the same physical template. See the set of photos of physical cylinder of defined geometry in Figure 5. The distance of the virtual camera is set at 0.6 m from the orthogonal plane where the first cylinder appears. Assuming a full-frame sensor of height 24 mm, we set the camera focal distance to 45 mm, which corresponds to human vision. This translates to a vertical viewing angle of $\approx 30^{\circ}$. The diameter of all cylinders is 59 mm, and the height of material on its surface is 120 mm; their viewing distances range from 0.6 to 4.4 m. A side-by-side comparison of real 59-mm wide cylinder with its rendering observed from distance 0.6 m is shown in inset at the top-right corner of Figure 4.

3.3. Experiment A—controlled, static stimuli

The task in our experiment was to identify the critical viewing distance for all analysed materials where their BTF rendering becomes indistinguishable from BRDF rendering. For this purpose, we used

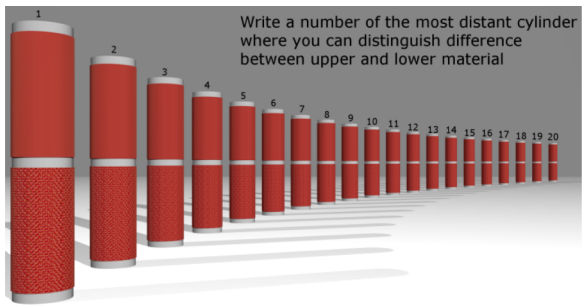

Figure 6: *Example of stimuli of controlled Experiment A—20 cylinders ranging in distance from 0.6 to 4.4 m (step 0.2 m).*

a test stimuli of 20 receding cylinders (example stimulus image in Figure 6) simulating distance 0.6–4.4 m, i.e. distance step between cylinders was 20 cm. The upper part of the cylinder always displayed the material's BRDF obtained by spatial averaging of its BTF, while the lower part always displayed material's BTF. The material renderings were pre-computed using ray-tracing (256 samples/pixel) with texture mip-mapping enabled. The subjects' task was: *Enter the number of the most distant cylinder where you can still distinguish difference between upper and lower material*. Thus the experiment type can be denoted as 20 alternative-forced-choice (AFC). Fourteen volunteer observers participated in the experiment. All had normal or corrected to normal vision and were naive with respect to the purpose of the experiment. Subjects evaluated 25 stimuli images, one per material. The resolution of images was 1920×1080 pixels and no resampling was applied. There was no time constraint to finish the task, and a typical session took about 5 min. All stimuli were presented on a calibrated 27" ASUS LCD display VG27AH (60 Hz, resolution 1920×1080 pixels), as the pixel-size was 0.311 mm the final displayed DPI was 82. The experiment was performed under dim room lighting and participants were advised to view the screen at a distance of 0.6m conforming to the designed scene geometry.

3.4. Experiment B—web-based, static stimuli

To validate results from the controlled experiment and to generalize its conclusion for different displaying scenarios, we performed a second psychophysical study, using a web-based testing interface. As the physical size of the scene has to be the same regardless of the size and DPI of a user's display device, we asked subjects to measure the width of a 500-pixel wide rectangle on the screen. This information allowed us to compute display DPI and resize stimuli images to an appropriate size. To avoid image blur due to images downsampling on the observer's device, the images were pre-computed in higher resolution 2160×1868 pixels. As the size of user display devices is often limited, we restricted the number of cylinders in the scene to 12; however, they span almost the same viewing distances (0.6–4.45 m) as distance between neighbouring cylinders is 35 cm. Therefore, the experiment type can be denoted as 12AFC. This step allowed us to reduce a physical width of displayed stimuli to images around 29 cm, which is the width of a typical 13" screen (see the example stimulus image in Figure 7).

Figure 7: *Example of stimuli of web-based Experiment B—12 cylinders ranging in distance from 0.6 to 4.45 m (step 0.35 m).*

The meaning of the stimuli images and task for the subjects was the same as in the *Experiment A*. The material renderings were precomputed using the OpenGL rendering (256 samples/pixel) with a texture mip-mapping enabled. We applied a high-quality filtering to suppress aliasing artefacts in BTF rendering. Forty volunteer observers participated in the experiment. Subjects evaluated 37 stimuli images, one per each material plus two additional variants for five anisotropic materials (material rotated for 45° and 90°) and two chessboard patterns. The subjects were strongly advised to observe the display from the distance of an outstretched arm (≈0.6 m) and to perform the experiment under dim lighting. The recorded average viewing parameters across all subjects are following: display resolution 1627 × 998 pixels, pixel size 0.279 mm, DPI 96.7. All stimuli images from the Experiment B are shown in Supporting Information.

3.5. Experiment C—controlled, dynamic stimuli

Finally, we replicated 12AFC *Experiment B* in a controlled environment as defined in *Experiment A*. The only difference was that instead of a static scene we used rotating cylinders. The task for the subject was the same as in the previous experiments, and stimuli images were computed using OpenGL. Eight volunteer subjects participated in the experiment, and the majority of them participated in *Experiment A* too. Subjects evaluated 25 stimuli images, one per each material. Movies were stored uncompressed with 30 frames/s. To avoid visually distractive seams near the material overlap on the cylinder, that could produce artificial lower frequency content, we placed a patch over this seam along the whole cylinder (as shown in the movie S1).

4. Experimental Results Analysis

This section describes the post-processing of the psychophysical data, presents an analysis of typical responses and discuses results of the psychophysical scaling.

4.1. Response filtering

The main limitation of our experiments using a virtual scene instead of a real scene lie in the limited DPI of a display device. This may potentially limit subjects in observing differences of materials having too low size of the smallest structure element with regards to the DPI, especially on high-DPI displays. Therefore, we resorted to the filtering of the subjects' responses. First, we identified the most distant cylinder where material structure elements were displayed on more than one pixel. The subject's response was recorded only if the selected number of cylinder was lower than the number of the identified cylinder. In other words, this guarantees that the subject still had a chance to see the cylinder where more than one pixel reproduced the biggest material structure size although she selected the one a step closer cylinder. This filtering removed 9.7% (34 of 350) responses in the *Experiment A*, and 8.5% (85 of 1000) responses in the *Experiment B*. The filtering took effect especially for materials with very small size of the structure element, namely, materials *fabric075*, *fabric106* and *leather16*. Despite the filtering, we obtained always more than 30% of available responses for each material.

Furthermore, we checked the user's vision by artificial chessboard patterns of size 2 and 5 mm. The visual angle of 2-mm square on the most distant cylinder subtended 0.026° (for DPI 100), which is still below a naked eye, typically of resolution 0.7 arc minute = 0.012° [Pir67]. In our experiment, over 90% of subjects were able to distinguish differences even on the last cylinder. Finally, we removed the subjects who were not able to spot a 5-mm pattern on the most distant cylinder as they probably either suffer from impaired vision or did not understand the task in the study (one subject). Also users who entered the measure of the calibration rectangle incorrectly in centimetres, and thus obtained an incorrect size of the stimuli images, were excluded from the study (two subjects).

Table 1: *Pearson correlations between results of the experiments.*

Correlation	ρ	p-Val
ExpA versus ExpB	0.966	0.000000
ExpA versus ExpC	0.811	0.000001
ExpB versus ExpC	0.889	0.000000

4.2. Averaged responses

The first step in the analysis of the perceptual data was computing the average critical viewing distance across all subjects. Figure 8 shows such data obtained from all three experiments across all 25 tested materials, divided into seven groups as described in Section 3.4. The errorbars represent standard deviation values computed across all subjects responses. Although we observe higher standard deviation values in data from the Experiment B (which was expected due to its uncontrolled nature), there is an apparent correlation between results of all experiments as shown in Table 1. The high correlation of the Experiment B with the other two validate results of this web-based study.

Along with responses, we also recorded the subjects observing time for individual stimuli images. The average time span was between 8 and 20 s, while the longest ones belonged to materials having very fine structure and thus were more difficult to distinguish.

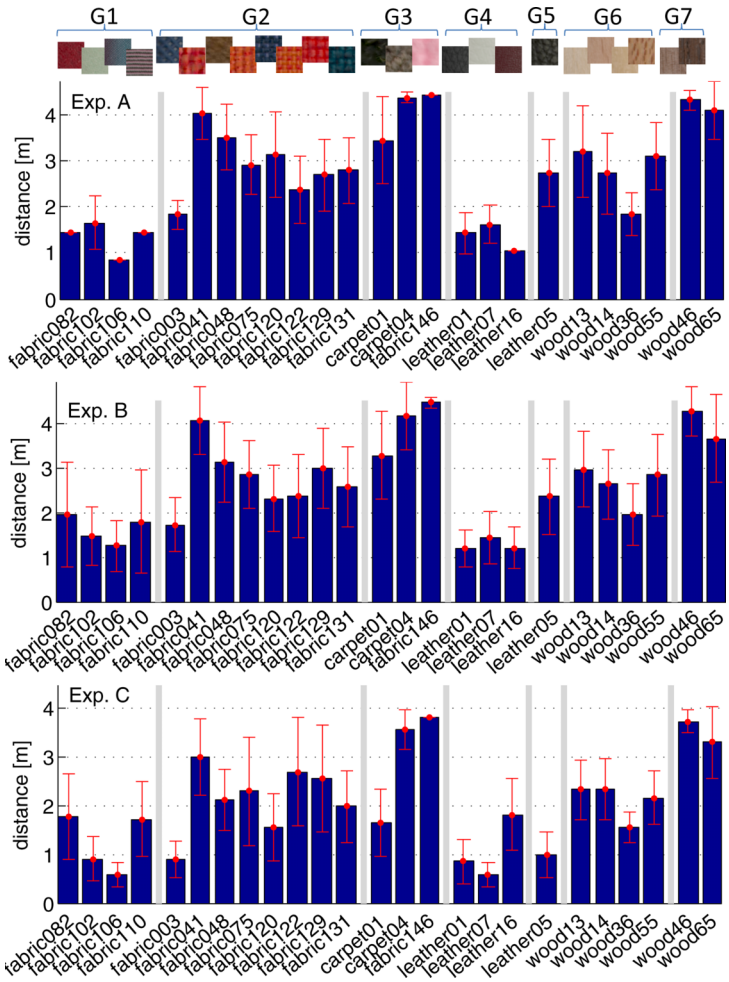

Figure 8: *Perceptually estimated critical distances, where on average the subjects were still able to distinguish between a BRDF and BTF. Included are results of all three experiments with standard deviations across all subjects.*

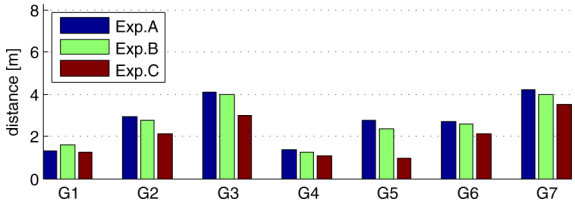

Figure 9: *The distances averaged across individual groups obtained in the three psychophysical studies (A, B and C).*

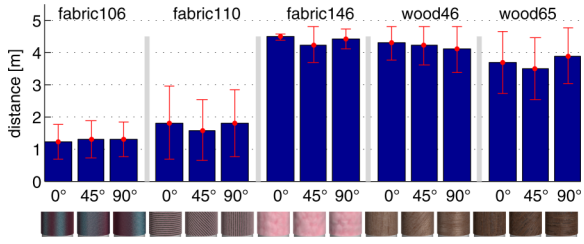

Figure 10: *Impact of anisotropic behaviour on distance where structure becomes indistinguishable.*

Surprisingly, longer observation times were also recorded for wood materials (except *wood46*), where the structure was much more apparent but missing sufficient contrast. The shortest times were recorded for rough materials where most users quickly spotted differences across all cylinders, e.g. *fabric146* or *wood65*.

Figure 9 depicts the estimated distances averaged across individual material categories. We observe that in the dynamic scene (Exp. C) the material structure became perceptually indistinguishable at closer distances (0.2–1.5 m closer) generally for all categories. The biggest difference was in groups having rough surface structure (G3,G5,G7). We assume that the reason for this behaviour may be in visual blurring of high-frequency features, present in these materials, due to material motion.

Finally, we analysed the estimated distance as a function of material anisotropy (*Experiment B*). The results for three fabric and

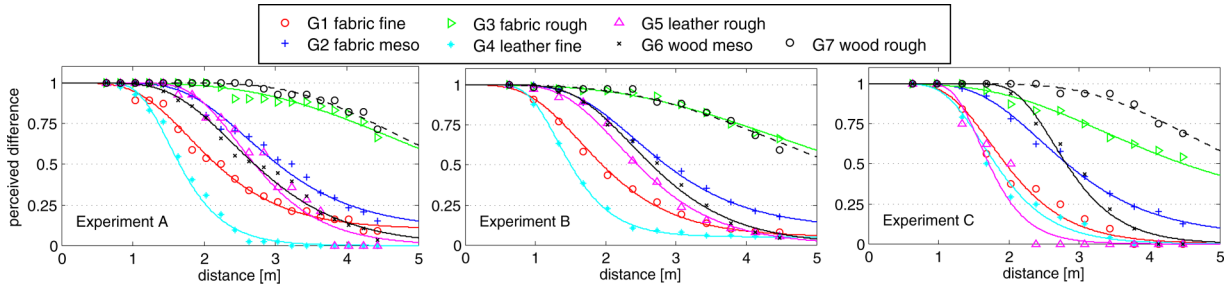

Figure 11: *Psychometric functions fitting data obtained from all three psychophysical experiments (A, B and C).*

two wood materials each having three different orientations over the cylinder (0°, 45°, 90°) are given in Figure 10. Differences within individual materials are below 0.5 m, which is presumably due to constant size, and thus visibility, of material structure regardless of changes in overall appearance. A more detailed analysis of anisotropy is shown in the Supporting Information.

4.3. Psychophysical scaling

For further analysis of subjects sensitivity in material structure recognition, we fitted data of each material category by a psychometric function [WH01]. For each observer, we recorded value 1 for those cylinders where the difference in structure was spotted, and 0 otherwise. The summarized data across all subjects was fitted by a Gumbel function (logarithmic Weibull function) using Palamedes psychophysical data analysis toolbox for MATLAB (http://www.palamedestoolbox.org/). Figure 11 shows fitted functions for data in individual material categories from all experiments. We can observe close similarity of the results from the static stimuli experiments (A and B).

In accordance with Figure 9, the perceived structure disappears first for material category *leather smooth* (G4), although a typical size of leather structure elements is much larger when compared with smooth fabric. This can be explained by their low surface profile and thus limited contrast. Then follows the category *fabric smooth* (G1) although its slope is considerably lower than from leather, which is presumably due to higher variability of material within this group. Similarly shaped and positioned psychometric functions were obtained for groups *fabric meso* (G2), *leather rough* (G5) and *wood meso* (G6). Similarly, closely resembling functions were achieved for groups *fabric rough* (G3) and *wood rough* (G7) marking completely different perceived distances.

The psychometric data from the dynamic stimuli experiment C show slightly different behaviour. Generally, the slants of the functions are steeper, which suggests a lower variance in subjects responses. This is supported by lower standard deviation values in Exp. C in Figure 8. The main difference is in the category of *leather rough*, where its structure becomes significantly less apparent. Similarly, the structure of smooth wood materials (G6) is also less apparent. This behaviour conforms with conclusions of Jarabo *et al.* [JWD*14] referring that for dynamic scenes the blur from motion results to higher visual equivalence than in static scenes. On the other hand, for *leather fine* category (G4) motion helps to recognize structure from a longer distance. We assume that this atypical behaviour is caused by the structure flickering of fine visual features in the material structure that cannot be easily detected in static scenes. This supports discussion with the subjects after the experiment, who mentioned that they were not comparing a BRDF with BTF, but rather focused on the presence of motion in the bottom part of a cylinder surface.

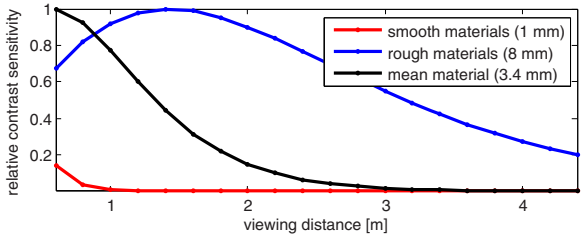

Figure 12: *A relative contrast sensitivity as a function of viewing distances in our test scene for the materials with smallest, biggest and typical size of structure element.*

Finally, the estimated critical distance thresholds for all three experiments and individual material groups for significance level 25% and 50% are shown in Figure 13. The level 25% means that only 25% or less of observers spotted difference between materials on a cylinder at given distance. While the level 50% should provide values of a typical user, we consider the level 25% as a safe adjustment for applied rendering algorithms.

4.4. Relationship to contrast sensitivity

As our work closely relates to the contrast sensitivity of the HVS, we implemented an achromatic CSF model of Burbeck and Kelly [BK80]

$$CSF(f_s, f_t) = 4\pi^2 f_s f_t \cdot e^{\frac{-4\pi(f_t+2f_s)}{45.9}} \cdot \left(6.1 + 7.3\left|\log_{10}\frac{f_t}{3f_s}\right|^3\right) \quad (1)$$

and plotted it as a function of viewing distance present in our test scene (see Figure 6). A temporal frequency f_t was set to 1 and spatial frequency f_s in cycles-per-degree was obtained from the size of material structure element. The result is shown in Figure 12. The CSF was evaluated for extreme and typical structure element sizes

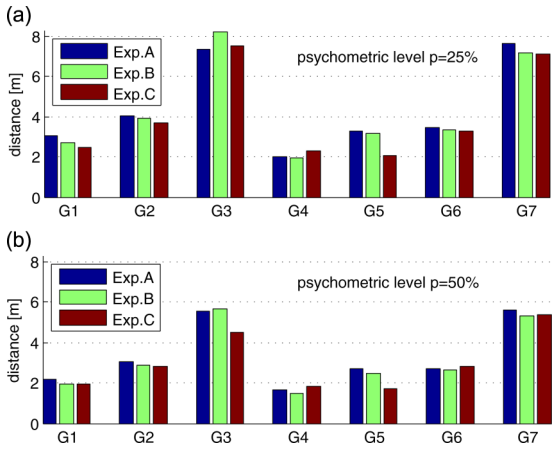

Figure 13: *The distances obtained by psychometric functions thresholding at levels (a) 25% and (b) 50%. Results are shown for individual material categories and all three psychophysical studies (A, B and C).*

Table 2: *Pearson correlations between results of the Experiment B and the tested statistical features (values over 0.8 are shown in bold).*

	Feature	ρ	p-Val
1	ExpB versus luminace	0.067	0.749150
2	ExpB versus contrast	0.583	0.002248
3	ExpB versus 1-SSIM	0.564	0.003326
4	ExpB versus 1-VDP2/100	0.649	0.000452
5	ExpB versus lum. contrast	0.616	0.001037
6	ExpB versus tex. contrast	0.506	0.009898
7	ExpB versus saliency	0.591	0.001875
8	ExpB versus structure size (S)	0.677	0.000203
9	ExpB versus S·luminace	0.499	0.011103
10	ExpB versus S·contrast	0.777	0.000005
11	ExpB versus S·(1-SSIM)	**0.878**	0.000000
12	ExpB versus S·(1-VDP2/100)	**0.808**	0.000001
13	ExpB versus S·lum. contrast	**0.941**	0.000000
14	ExpB versus S·tex. contrast	**0.934**	0.000000
15	**ExpB versus S·saliency**	**0.943**	0.000000

present in our materials (see Figure 3), i.e. 1, 8 and 3.4 mm. When compared to the psychometric functions in Figure 11 it is apparent that although the CSF itself can predict HVS sensitivity to frequencies of typical materials relatively correctly, it is unable to predict reliably the sensitivity to materials of extreme sizes, where it tends to underestimate the critical distance. This is probably due to lower texture contrast in the real materials appearance, that is different from idealized sinusoidal-wave gratings assumed by the CSF model as well as due to impact of material illumination. Therefore, we focused on texture contrast features in our model of critical distance. Moreover, the tested spatiotemporal CSF model shows that HVS contrast sensitivity decreases with increasing temporal frequency, and thus guarantees the visual temporal coherence of renderings in applications relying on the estimated critical distances.

4.5. Limitations

Note that the type of motion in *Experiment C* is quite rare, as in many other scenes object's geometry combined with movement would presumably visually mask these fine visual differences and thus increase the critical viewing distance.

Due to a limited display DPI, we filtered a certain portion of subjects responses (Section 4.4) that correspond to more distant cylinders. This unfortunately slightly bias the critical distance of materials with very fine structure (*fabric082*, *fabric106*, *fabric110*, *leather16*) towards ≈0.2 m lower distances. This filtering effect for these materials can be safely compensated in applications intended for standard screens with DPI around 100, where the users would face the same DPI limitation as subjects in the experiment.

When we compare distances derived from psychometric scaling Figure 13 with simple distances averaging in Figure 9, we obtain higher perceived distances for more rough surfaces for the psychophysically scaled variant. This is due to a limited span of distances available in the experiment resulting in saturated subjects responses for very rough materials, i.e. subjects were forced to select the last cylinder although they might be able to spot difference even further. Unfortunately, we could not increase viewing distance due to a limited resolution of screen, so a proper analysis of rough materials structure visibility (structure element size above 10 mm) becomes a subject of future endeavours.

5. Structure Visibility Prediction

This section compares obtained psychophysical results with the combination of standard image and computational metrics. The intention is to understand clues behind the visual perception of material structure. To this end, we analysed the correlation of psychophysical results with several promising features. We used data from *Experiment B* due to the highest number of subjects (40) as our golden standard.

The tested computational features were: (1) average luminance of the cylinder, (2) contrast approximated by mean standard deviation of difference between BTF and BRDF, (3) $1 - \text{SSIM}$ (Structure Similarity Index) [WBSS04] between BTF and BRDF, (4) $1 - \text{VDP2}/100$ (Visual Difference Predictor) [MKRH11] between BTF and BRDF, (5) luminance contrast C_L, (6) texture contrast C_T and (7) saliency obtained as linear combination of the last two features $C_L + 10 \cdot C_T$ [PN04]. We have also tested features based on luminance and chromatic contrast estimated from pixelwise cone responses [WM97], however, their performance was relatively poor. All the selected features were computed from the greyscale bottom cylinder area of stimuli images from *Experiment B*, and values of most of them were obtained by the averaging of values across the image plane.

The correlation between the experiment and features 1-7 are shown in the first part of Table 2. The obtained values are not very high, which suggests that there is another element perceived by subjects not fully captured using the tested features. As an important factor impacting the recognition of material structure is its scale, we tested as an additional feature (8) the largest size of

Table 3: *Estimated linear parameters fitting three selected features to the data from Experiment B.*

Feature	k_1 (std)	k_2 (std)
S·(1-SSIM)	1.582 (0.044)	1.748 (0.053)
S·(1-VDP2/100)	1.777 (0.033)	18.694 (0.384)
S·saliency	1.292 (0.023)	1.201 (0.015)

material structure element (see Figure 3). As expected, this gave the best correlation value so far. Therefore, as another step we used structure size for the scaling of results of individual tested features and computed correlation with the experiment again. These values, shown at rows 9–12 in second part of Table 2, show a very strong correlation especially for scaled variants of SSIM, VDP2 and image saliency.

While these features seem to be very promising candidates for critical viewing distance prediction, they are just proportional to experimental data and have very low values for materials where computed difference (SSIM, VDP2) or contrast are too low. Therefore, we suggest the adaptation of the features values F to the result of experiment E using a simple linear model $E = k_1 + k_2 F$. We used robustfit function in MATLAB for parameters computation. To check the stability of the parameters, we performed a leave-one-out parameters estimation, i.e. distance value for particular material was obtained from parameters computed on values of all-but-this material. Obtained standard deviation values of such 25 sets of parameters are very low. The mean k_1, k_2 values and their standard deviations are shown in Table 3. This analysis revealed that the most stable parameters were achieved for image saliency feature. Another advantage of this feature is that, due to its Gaussian pyramid, it does not rely so much on the selection of a structure neighbourhood size (SSIM) or display information and observer distance (VDP2). Typical times needed for evaluation of individual features on our stimuli images were: SSIM 0.005 s, VDP2 0.270 s and saliency 0.082 s.

Parameters of individual tested methods:

SSIM	$C_1 = 0.01$, $C_2 = 0.03$, Gaussian window size 11 pixels, $\sigma = 1.5$
VDP2	Screen size: 24", resolution 1900 × 1200, viewing dist. 0.6 m
Saliency*	Gaussian window size 11 pixels, $\sigma_c = 0.5$, $\sigma_s = 2.5$, $\beta = 2$, $s = 5$

*Our implementation in MATLAB is available at http://staff.utia.cas.cz/filip/projects/15CGF

Finally, the predicted distances are shown in Figure 14 together with errors, where red bars show the proportion of underestimation of the correct critical distance, while the green ones its overestimation. The best performance was achieved by the image saliency feature (c). When compared with distances obtained from *Experiment B* (Figure 8), it is apparent that errors of our model are within standard deviations across subject responses. This supports our conclusion that the proposed image saliency feature [PN04] scaled by material structure size provides a promising and computationally reasonable model for material structure visibility.

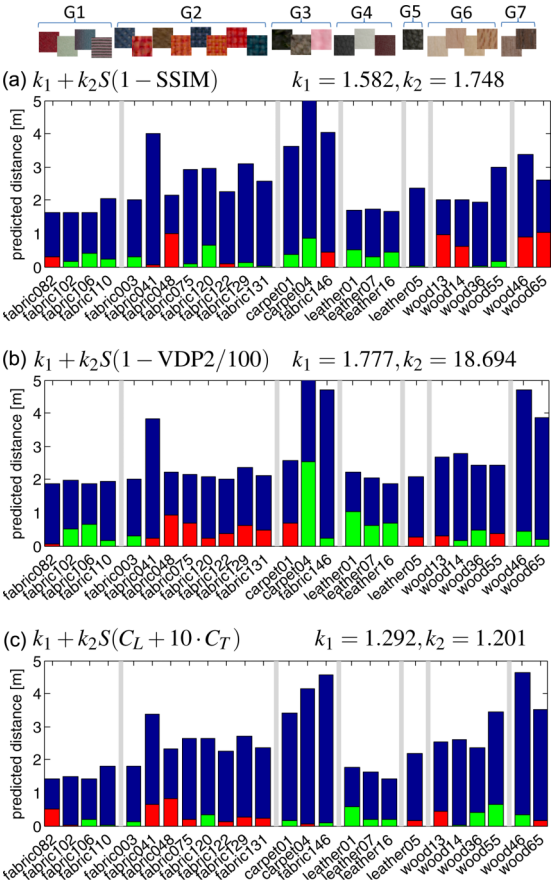

Figure 14: *Comparison of critical viewing distance prediction using structure sized features based on (a) SSIM, (b) VDP2, (c) image saliency. Red bars show the proportion of underestimation of the correct distance while the green ones its overestimation.*

Finally, Figure 15 compares values of the critical distances as obtained from the *Experiment B* (E) and from the proposed prediction (P) for a typical representatives of individual tested material categories.

6. Applications

Although a position of observer with respect to viewed objects in virtual reality is not constrained in general, we assume that there exist typical constrained viewing scenarios. For instance, a car interior is typically viewed and rendered from a driver's perspective. Similarly a room interior in virtual-walk-through applications is visualized from a perspective of its visitor usually located near entrance or centre of the room. In other words, we assume that our method can be beneficial in any applications where viewing distance constraints can be imposed or assumed. As a consequence, it allows to render a more complex scenes (with more materials) using the same HW.

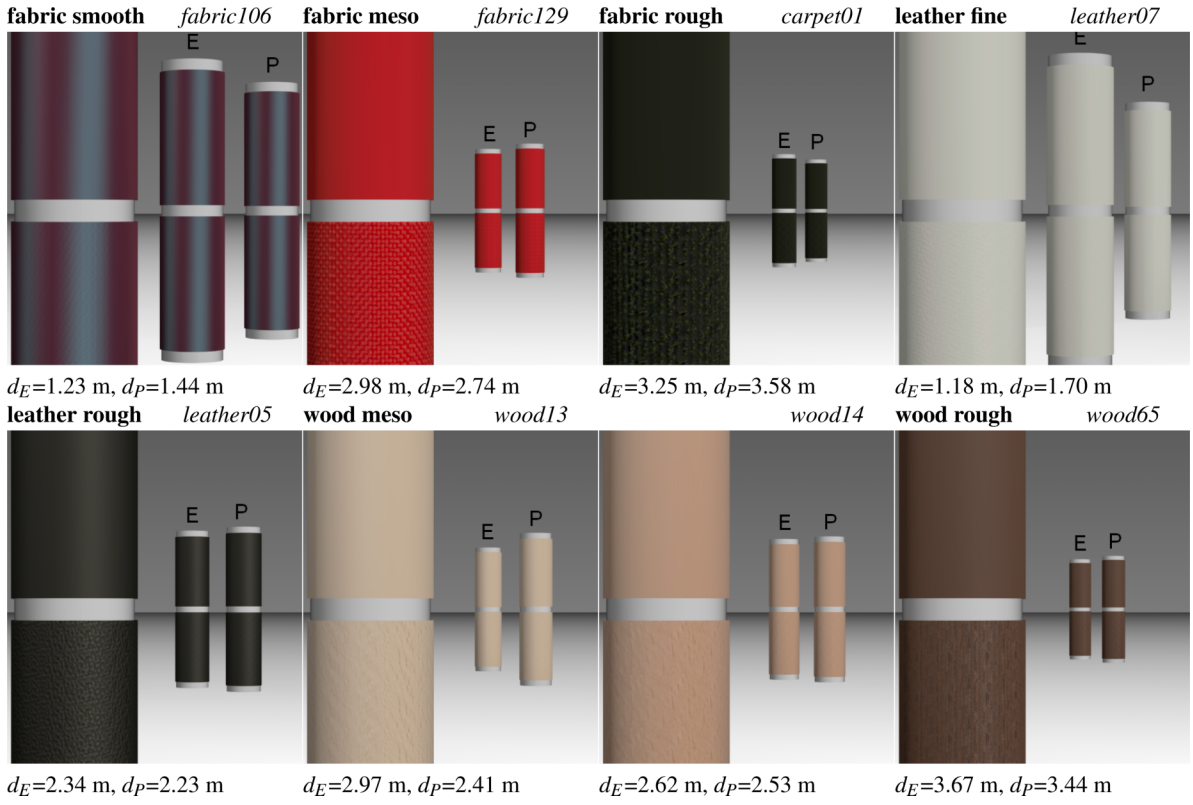

Figure 15: *A critical distance example comparing 0.6-m distant reference column (left) with columns corresponding to distances obtained by averaging across subjects responses from Experiment B (E) and from the statistical predictor (P). To conform with the real scene geometry, the image should be zoomed so as a width of the left reference column subtends 59 mm on the screen and the screen is assumed to be viewed from distance 0.6 m.*

Table 5: *Measured rendering speed (in frames-per-second) of scene 1920 × 1080 pixels with various material appearance representation approaches, and PSNR of both approaches related to the reference full BTF method.*

Method	BTF	BRDF	Proposed
FPS	8.3	20.2	11.8
PSNR [dB]	–	46.3	58.6

Even though we would assume a non-constrained viewing scenario where storing of full BTF in memory is necessary, our method can reduce computational cost related to BTF spatial interpolation when the viewing distance exceeds the critical one. Although this interpolation is often HW supported (depending on data structuring) its removal can save typically four BTF reconstructions (needed for spatial interpolation) and substitute them by a single cheaper BRDF reconstruction. Although one can claim that our results can be substituted simply by texture mip-mapping, we argue that our results are more restrictive as the estimated critical distance is in majority cases shorter than the one corresponding to the lowest mip-mapping level (i.e. mapping the entire material structure onto a single pixel).

To support our claims, we created a virtual scene consisting of chairs in a room. The chairs are organized in six rows, each consisting of 12 chairs. The upper cushioned parts of chairs were covered by fabric/leather materials while the bottom construction was covered by wood materials. The floor consists of two different carpets and walls are represented by wood (total 14 materials). Depth range of the room was 12 m (Figure 16 c) and materials were mapped on the object in a physical scale. The scene is illuminated by two point lights. First, we rendered the scene in full BTF representation as it is shown in Figure 16(a). Then we implemented an OpenGL shader that switches texel's rendering from BTF to BRDF where the viewing distance exceeds the predicted critical distance d_E for a given material. The BTFs in the scene were filtered to remove aliasing artefacts. The result of the shader is shown in Figure 16(b) and its difference to BTF only rendering scaled 30× is shown in Figure 16(d).

One can see that visual differences are negligible and Peak Signal-to-Noise Ratio (PSNR) between full BTF rendering and our approach rendering is 58.6 dB as shown in Table 5. The table also highlights the gain in computational speed. While full BTF rendering of the scene in resolution 1920×1080 pixels achieves 8.3 FPS, BRDF rendering (as our upper bound) 20.2 FPS, our method combining both approaches based on the predicted critical distance achieves 11.8 FPS. This represents 30% gain in speed when compared to full BTF rendering without significant loss of perceived fidelity. Although, these values will vary across

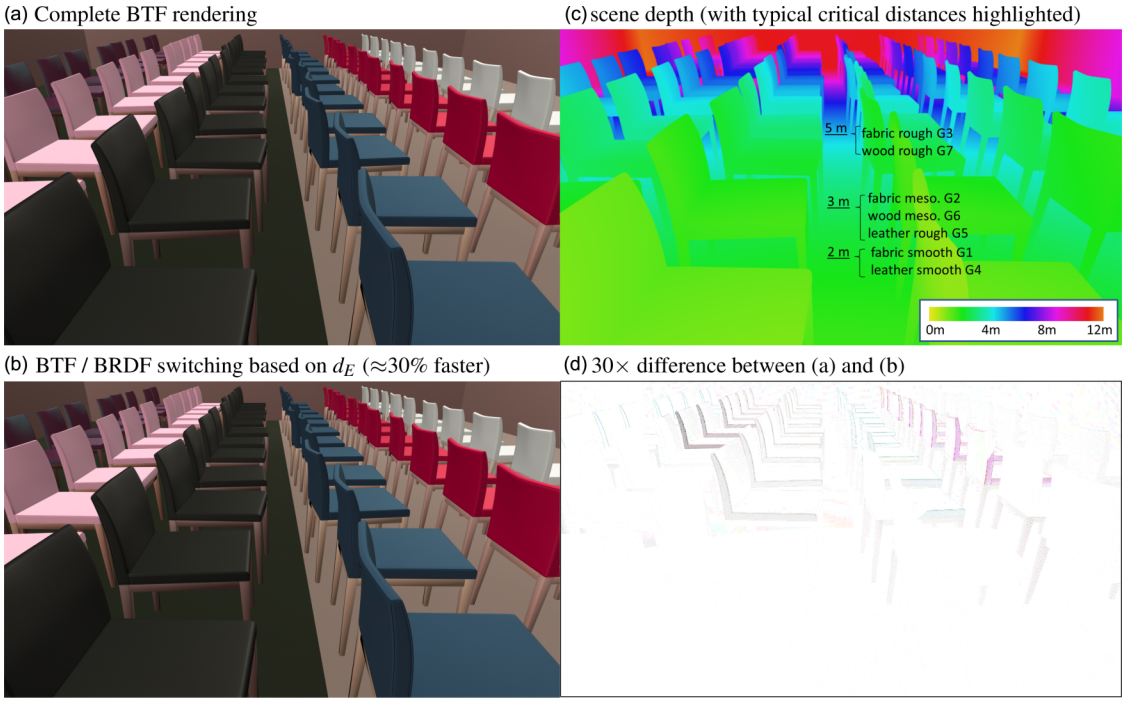

Figure 16: *Scene consisting of six rows of chairs and 14 materials: (a) rendering using BTF only, (b) switching bet\ween BTF and BRDF based on the proposed model of critical distance d_E, (c) scene range-map (depth 12 m), (d) 30× scaled difference between (a) and (b). The proposed approach achieves 30% faster rendering performance while maintain high visual fidelity to full BTF rendering (PSNR 58.6 dB).*

different scenes, we consider this experiment as an important showcase application of the proposed critical distance. We have not encountered any temporal discontinuities in our material appearance renderings (see the accompanying movies S1 and S2).

All tests were performed on a PC with Intel Core 5 2500 3.3 GHz, 16 GB RAM and graphics nVidia GeForce GTX 570. The relatively low speeds in Table 5 result from the high screen resolution (Full HD), where all pixels are evaluated. Furthermore, as we are not using an analytical BRDF model, we have to interpolate the BRDF measurements (from nine values using barycentric interpolation in each spectral channel). Finally, we use just unoptimized shader and data are reconstructed for each light independently, which also prolongs the rendering times.

Another application of our findings can be an instrumental tool that would help interior designers to select materials according to their distance and intended appearance in the virtual environment. The critical distance can create a map of the scene (similar to Figure 16 c), e.g. highlighting objects where the fabric structure would be visible from a given viewpoint.

7. Conclusions

The main objective of this paper was the analysis of material structure visibility in virtual environments as a function of viewing distance. Identification of the so-called *critical viewing distance*, i.e. material-dependent viewing distance where material structure/texture can be visually distinguished, is important in the decision of a BRDF replacing a BTF to avoid a loss of realism. On the other hand, an application of the critical distance can save computational costs in situations where a BTF may be substituted by a BRDF without the loss of visual fidelity. To this end, we estimated such distances for 25 interior materials in two static and one dynamic psychophysical studies with up to 40 participants. We divided our materials into seven categories and identified the critical viewing distances for each of them. We found that critical distances range from 2 m for smooth fabric and leather through 4 m for rough leather, moderately structured fabric, smooth wood to 8 m for rough fabric, carpet and wood. Moreover, we analysed a number of computational features to derive a reliable model of psychophysical results applicable as a predictor of critical distances for new materials. The best modelling performance was obtained by a feature based on the image salience scaled by the size of the largest material structure element.

Our results can benefit a number of applications dealing with complex virtual scene renderings. We have shown that the application of a critical viewing distance can reduce rendering costs significantly. Furthermore, it can serve as a useful tool for interior designers in selecting of materials according to an intended appearance in virtual environment. In the future, we plan to analyse the impact of illumination in an environment on the critical viewing distance.

Acknowledgements

We thank the anonymous reviewers for their insightful and inspiring comments. Also we would like to thank all volunteers for taking part

in the psychophysical experiments. This research has been supported by the Czech Science Foundation grant 14-02652S.

References

[BK80] Burbeck C. A., Kelly D.: Spatiotemporal characteristics of visual mechanisms: Excitatory-inhibitory model. *Journal of the Optical Society of America 70*, 9 (1980), 1121–1126.

[DvGNK99] Dana K., van Ginneken B., Nayar S., Koenderink J.: Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics 18*, 1 (1999), 1–34.

[FCGH08] Filip J., Chantler M., Green P., Haindl M.: A psychophysically validated metric for bidirectional texture data reduction. *ACM Transactions on Graphics 27*, 5 (December 2008), pp. 138:1–138:11.

[FH09] Filip J., Haindl M.: Bidirectional texture function modeling: A state of the art survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence 31*, 11 (October 2009), 1921–1940.

[FHC10] Filip J., Haindl M., Chantler M.: Gaze-motivated compression of illumination and view dependent textures. In *Proceedings of the 20th International Conference on Pattern Recognition (ICPR)* (August 2010), doi: 10.1109/ICPR.2010.217, pp. 862–864.

[GMSK09] Guthe M., Müller G., Schneider M., Klein R.: BTF-CIELab: A perceptual difference measure for quality assessment and compression of BTFs. *Computer Graphics Forum 28*, 1 (2009), 101–113.

[JWD*14] Jarabo A., Wu H., Dorsey J., Rushmeier H., Gutierrez D.: Effects of approximate filtering on the appearance of bidirectional texture functions. *IEEE Transactions on Visualization and Computer Graphics 20*, 6 (June 2014), 880–892.

[KFB10] Křivánek J., Ferwerda J., Bala K.: Effects of global illumination approximations on material appearance. *ACM Transactions on Graphics 29*, 4 (2010), pp. 112:1–112:10.

[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 Transactions on Graphics 30*, 4 (2011), 40:1–40:14.

[MMK*06] Meseth J., Müller G., Klein R., Röder F., Arnold M.: Verification of rendering quality from measured BTFs. In *Proceedings of Third Symposium on Applied Perception in Graphics and Visualization (APGV)* (July 2006), doi: 10.1145/1140491.1140517, vol. *153*, pp. 127–134.

[NRH*77] Nicodemus F., Richmond J., Hsia J., Ginsburg I., Limperis T.: *Geometrical Considerations and Nomenclature for Reflectance*. NBS Monograph 160, National Bureau of Standards, U.S. Department of Commerce (1977), pp. 1–52.

[Pir67] Pirenne M.: *Vision and the Eye*. Chapman and Hall, London, 1967.

[PN04] Parkhurst D. J., Niebur E.: Texture contrast attracts overt visual attention in natural scenes. *European Journal of Neuroscience 19*, 3 (2004), 783–789.

[QKM*00] Qi J., Kerr Y., Moran M., Weltz M., Huete A., Sorooshian S., Bryant R.: Leaf area index estimates using remotely sensed data and BRDF models in a semiarid region. *Remote Sensing of Environment 73*, 1 (2000), 18–30.

[RFWB07] Ramanarayanan G., Ferwerda J., Walter B., Bala K.: Visual equivalence: Towards a new standard for image fidelity. *ACM Transactions on Graphics 26*, 3 (2007).

[SGS*02] Schaaf C. B., Gao F., Strahler A. H., Lucht W., Li X., Tsang T., Strugnell N. C., Zhang X., Jin Y., Muller J.-P., Lewis P., Barnsley M., Hobson P., Disney M., Roberts G., Dunderdale M., Doll C., d'Entremont R. P., Hu B., Liang S., Privette J. L., Roy D.: First operational BRDF, albedo nadir reflectance products from MODIS. *Remote Sensing of Environment 83*, 1 (2002), 135–148.

[SH05] Somol P., Haindl M.: Novel path search algorithm for image stitching and advanced texture tiling. In *Proceedings of WSCG* (2005), pp. 155–218, http://wscg.zcu.cz/wscg2005/Papers_2005/Full/!WSCG2005_Full_Proceedings_Final.pdf.

[SSK03] Sattler M., Sarlette R., Klein R.: Efficient and realistic visualization of cloth. In *Proceedings of Eurographics Symposium on Rendering* (2003), pp. 167–178, http://cg.cs.uni-bonn.de/en/publications/paper-details/sattler-2003-efficient/.

[VLD07] Vangorp P., Laurijssen J., Dutre P.: The influence of shape on the perception of material reflectance. *ACM Transactions on Graphics 26*, 3 (2007).

[Wan95] Wandell B. A.: *Foundations of Vision*. Sinauer Associates, Inc., Sunderland, MA, 1995.

[WBSS04] Wang Z., Bovik A., Sheikh H., Simoncelli E.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing 13*, 4 (2004), 600–612.

[WH01] Wichmann F., Hill N.: The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics 63*, 8 (2001), 1293–1313.

[WM97] Webster M. A., Mollon J.: Adaptation and the color statistics of natural images. *Vision Research 37*, 23 (1997), 3283–3298.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web site:

Online Appendix.

Video S1.

Video S2.