EG UK Computer Graphics & Visual Computing (2016)
Cagatay Turkay and Tao R. Wan (Editors)

Selective BRDFs for High Fidelity Rendering

T. Bradley[1,2], K. Debattista[1], T. Bashford-Rogers[1], C. Harvey[1], E. Doukakis[1] and A. Chalmers[1]

[1]University of Warwick, UK
[2]t.e.bradley@warwick.ac.uk

Abstract

High fidelity rendering systems rely on accurate material representations to produce a realistic visual appearance. However, these accurate models can be slow to evaluate. This work presents an approach for approximating these high accuracy reflectance models with faster, less complicated functions in regions of an image which possess low visual importance. A subjective rating experiment was conducted in which thirty participants were asked to assess the similarity of scenes rendered with low quality reflectance models, a high quality data-driven model and saliency based hybrids of those images. In two out of the three scenes that were evaluated significant differences were not found between the hybrid and reference images. This implies that in less visually salient regions of an image computational gains can be achieved by approximating computationally expensive materials with simpler analytic models.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Rendering—Reflectance Modelling

1. Introduction

There is a continuing demand for increased accuracy and simulation speed in virtual environments. This comes at the cost of ever increasing demands on computational resources, especially when using algorithms relying on ray-traced lighting, and in scenes with multiple complicated materials. These demands can be partially mitigated through various strategies, such as improving light transport algorithms, coherence for tracing and shading, filtering, and reducing computation in less visually important regions of the scene. This paper investigates the latter approach, specifically whether straightforward and computationally inexpensive surface reflection models can be used in place of accurate but more detailed models in less salient regions of the image.

Visual attention models provide a framework to predict the areas of an image which are likely to be attended to by the Human Visual System (HVS). These models have been applied to improving rendering performance by several previous authors [CCW03] [CDMPdS07] [KDCM14] [GDS14], however this work has predominantly focused on decreasing the number of samples computed in a Monte Carlo image synthesis context. Our work aims to improve performance in an orthogonal manner; by reducing the complexity of materials in less visually salient areas, computational savings can be gained without a significant loss in perceived quality of the resultant image.

Material models used in physically-based rendering are known as *Bidirectional Reflectance Distribution Functions* (BRDFs). These are 4D descriptions of how incident light is reflected in a given direction. Two types of models are frequently used in rendering systems; analytical and data-driven. Analytical BRDFs range from simple models of diffuse reflectance, to parametrised glossy models, such as Phong [Pho75] and Walter [WMLT07], and multi-lobed BRDFs [LFTG97]. These parameters can be altered to affect the appearance of multiple surfaces, but can only approximate real materials. Data-driven BRDFs consist of captured reflectance values of real-world materials. These typically consist of a 4D (or 3D in the case of isotropic materials) lookup table which is queried at runtime. Uncompressed data-driven BRDFs require significant storage, for example the MERL database [MPBM03] stores each measured surface as a 33MB binary file, the total size of the database of 100 materials being 3.30 GB. This requires significant memory bandwidth during the frequent lookups from the table.

Typical rendering systems additionally use mixtures of these BRDFs, often in a spatially varying [Mca02], layered [WW07], or as a mixture driven by a shader. Evaluating shading on a surface can take longer than tracing rays [ENSB13], and therefore savings in BRDF evaluation can potentially significantly speed up the rendering process. The focus of this paper is on approximating data-driven BRDFs with analytical models in less salient regions. We choose data-driven BRDFs as a baseline as they are expensive to evaluate and provide accurate measured data.

In this paper, an application of visual saliency models to reduce the computational and memory bandwidth costs for BRDF evaluation is described. A series of analytic BRDF models were fitted to the MERL dataset, and applied to several virtual scenes. A saliency map was generated for each scene and used to evaluate whether

© 2016 The Author(s)
Eurographics Proceedings © 2016 The Eurographics Association.

users notice a difference between an expensive, but accurate, data-driven BRDF, and a fitted analytic approximation.

The structure of this paper is as follows: Section 2 discusses related work, Section 3 outlines the methodology that was used to prepare the scenes and reflectance models, as well as the procedure and design of the perceptual experiment. The results of the perceptual tests are given in Section 4, discussed in Section 5 and conclusions are drawn in Section 6.

2. Related Work

This section describes related work from the domains of applications of visual attention to accelerating rendering, as well as BRDF models.

2.1. Visual Attention and Level of Detail

Visual attention models attempt to measure Regions of Interest (ROIs) in an image, either from a *Top-Down* or *Bottom-up* perspective. Top-down models are task focused, and estimate ROIs based on an objective that the observer is trying to achieve. Bottom-up models estimate ROIs based on knowledge of the Human Visual System (HVS) and *preattentive features*. A good summary of preattentive features can be found in Healy and Enns' paper [HE12].

The saliency map was proposed by Koch and Ullman [KU87]. Saliency maps are grayscale images, where each pixel value corresponds to how likely an observer is to pay attention to that point in the scene, the higher the pixel value the higher the saliency of that point.

Itti and Koch [IKN98] then developed a computational model of bottom-up visual attention by using a scale-space feature detection method on an image, that is they subsampled and smoothed the image to emphasise prominent features or *conspicuities* within the image. Their method performed well on a variety of images but struggled in images with uniform noise. They also had no eye trackers or other methods available to validate their model by testing it with human subjects.

Cater *et al.* [CCL02] examined whether *inattentional blindness* could be used to reduce the quality of portions of a rendered image, in order to reduce computation time. Their experiment verified their model of task-based visual attention using eye tracking technology and revealed that users were not able to consistently notice that areas in the images that were of a reduced quality. Cater *et al.* [CCW03] expanded their work and combined task maps and a contrast sensitivity function to selectively reduce the number of samples in unimportant regions of an image and regions where errors would be more noticeable. This reduced the resolution of the image in less salient areas but occasionally stuttered when interpolating between neighbouring frames.

Sundstedt *et al.* [SCCD04] developed task importance maps, which allowed users to manually tag objects to be rendered at higher quality, which would be relevant to a given task. They determined that if sufficiently focused on a task observers would often fail to see reductions in image quality even if those low quality regions were within the foveal region. They later expanded this method [SDL*05] combining task maps with saliency maps to create overall importance maps. Harel *et al.* introduced the *Graph-Based Visual Saliency* (GBVS) model [HKP06], which used normalized *activation maps* and graph cuts to identify salient regions within an image, their method combines well with other methods. However their method favours the centre of an image. Longhurst *et al.* [LDC06] introduced a GPU-based system with live anti-aliasing within a selective rendering framework. Chalmers *et al.* [CDM-PdS07] created low quality *Snapshots* of scenes using rasterization and then subdivided the image into salient regions, sampling important sub-images more frequently to achieve a higher level of perceived realism.

More recently Koulieris *et al.* [KDCM14] developed a system based on top down visual saliency and context based object tagging to reduce the level of detail of subsurface light transport, refraction and bump mapping in unimportant regions.

2.2. BRDFs and Fitting

In order to reproduce the appearance of real materials digitally it is necessary to either have stored data, which represents its appearance, or determined a set of parameters for a BRDF, which will produce a function that closely matches that data. Finding parameters for representing a specific material is referred to as *fitting* that BRDF to the material. This section will focus on the literature related to acquisition of reflectance data and the fitting of BRDFs to that data.

Gonioreflectometers have been employed to capture how light bounces off real world materials. Dana *et al.* [DVGNK97] gathered sparse measurements from over 60 material samples, forming the *CUReT* database, and fitted them to the Oren-Nayar [ON95] and Koenderink [KVDS96] BRDFs.

The MERL material database [MPBM03] was the first large database of densely sampled data-driven isotropic materials, containing a variety of plastics, metals, fabrics and natural materials. It has been used as a benchmark for comparing the flexibility and accuracy of new BRDFs as a replacement for acquiring first-hand data [BSH12] [LKYU12] [BLPW14] [HP15].

Two of the BRDFs used in this study are the Phong BRDF [Pho75] and the Walter BRDF [WMLT07]. The Phong method simulates glossy reflections through the use of a cosine lobe raised to an exponent alongside a diffuse component.

The Walter method is an adaptation of the Cook-Torrance [CT82] model, designed to simulate reflection and refraction from rough surfaces through the use of the GGX microfacet distribution function and a correction to Smith's shadowing and masking function [Smi67]. This is based on a physically plausible model of the underlying surface, and has been used to represent a wide variety of materials.

Lafortune *et al.* [LFTG97] developed a system to fit multiple Phong lobes to measured materials, which proved more effective than individual lobes however they found this insufficient for some materials and found that the stability of their fitting function decreased as the number of reflectance lobes increased.

Ngan *et al.* [NDM05] used Sequential Quadratic Programming

© 2016 The Author(s)
Eurographics Proceedings © 2016 The Eurographics Association.

to minimise a squared difference error metric. They evaluated the ability of seven analytical BRDFs to represent the MERL database. The Cook-Torrance [CT82], Ashikhmin-Shirley [AS00] and He *et al.* [HTSG91] BRDFs fit the data well with minimal errors, while simpler Blinn-Phong [Bli77] and Lafortune [LFTG97] models performed poorly. The grazing angles were not included in the fitting, as the data was extrapolated when the MERL materials were captured. They also found that it was difficult to fit many of the materials with a single reflectance lobe. The disadvantage of using an extra reflectance lobe is the increased computation time and fitting becomes less stable.

Recently a number of BRDFs have been designed, specifically to fit to the MERL database, Bagher *et al.* [BSH12] introduced the Shifted Gamma microfacet Distribution (SGD) function for the Cook Torrance BRDF [CT82], replacing the traditional Beckmann distribution. Their reflectance model has a large number of parameters but can fit each material in the database with a single reflectance lobe. They precompute the values of their shadowing and masking function for offline rendering and use an approximation for GPU-based rendering as the precomputed values are slow to access in real-time applications.

Löw *et al.* [LKYU12] introduced the ABC BRDF which accurately models glossy surfaces. Brady *et al.* [BLPW14] proposed genBRDF which used genetic algorithms to generate new BRDFs from the MERL database. Holzschuch and Pacanowski's BRDF [HP15] designed a physically based BRDF incorporating both reflectance and diffraction, which provided accurate fits to the MERL database and outperformed the SGD distribution for the Cook-Torrance BRDF, which had previously provided the best fits to the MERL database.

3. Methodology

This work is motivated by the need to provide better overall performance for scenes that may not require detailed material representations at every point in the scenario. In order to demonstrate the feasibility of such a system an experiment was conducted to identify perceptual differences amongst analytical and data-driven BRDFs and hybrids of the two that use the higher quality BRDF in areas of the region considered more salient.

3.1. Design

The experiment is a subjective rating study, in which participants were asked to rate the quality of images in comparison to a ground truth image on a scale of 1 to 100. A hidden reference is also included to provide a relation to the ground truth image. The rating design permits quantification of distance between methods. The hidden reference enables comparisons with the other stimuli to identify perceivable differences across them.

The scale used in this experiment offers a sufficient breadth for participants to give a wide range of ratings. Each participants' scores are normalised when calculating results in order for effective comparisons to be made.

The independent variables are the analytical BRDF which is used to create the mixed images and the scenes used. Both independent variables follow a within-participant design. The BRDF variable consists of seven possibilities, three analytical BRDFs, a diffuse model (D), the Phong BRDF (P) [Pho75] and the Walter BRDF (W) [WMLT07], three mixed saliency models consisting of the three chosen BRDFs (SD, SP, SW respectively) mixed with the data-driven BRDF using a saliency map and a hidden reference (R). The three scenes were of a kitchen, a conference room, and a lounge, see Figure 2.

The camera angle, BRDF fitting procedure, sampling algorithm, saliency model, image resolution, number of samples per pixel and viewing time for each image were set as constant across all scenes and reflectance models. The dependent variable is the rating given to each stimulus.

3.2. Materials

This section describes the preparation of the materials, particularly the stimuli used in the experiment.

3.2.1. BRDFs

The choice of the data driven BRDF used in these experiments is primarily motivated by three factors, the number of available materials, the density of the measurements and the focus of related literature. Therefore the MERL database was chosen as the reference BRDF. The Walter BRDF was chosen because it is heavily used in industry, fits the MERL database well and is a good representation of micro-facet BRDFs. The Phong BRDF was also chosen for its prevalence in rendering and its speed. It is purely specular and is normally combined with a basic lambertian diffuse model to provide colour.

For the purpose of this experiment a densely populated data driven BRDF (the MERL database) was chosen as a reference, as using an analytic model would require an extra level of fitting. For this purpose Löw's [LKYU12] fitting function is adapted to fit to the methods that are used in the perceptual experiment. It was chosen for its ability to fit to simple models rather than to the ABC method for which it was designed.

This section outlines the procedure that was followed to fit the analytical BRDFs used in this study to each entry in the MERL database. Each model had independent parameters for each colour channel, with the exception of roughness parameters, as surface micro-structure is independent of the incident wavelength. In the case of the Walter BRDF, the incident index of refraction was the same for each colour channel as the incident medium is air and the difference in refractive index for red, green and blue's respective wavelengths is negligible.

The MERL BRDFs have an angular resolution of one degree, as the materials are isotropic, the reflectance doesn't vary with the incident angle, ϕ, therefore $\phi = -\phi$. The density of the tabulated BRDFs is then $90 \times 90 \times 180 = 1,458,000$ entries for each colour channel.

In order to find parameters for the diffuse, Phong and Walter BRDFs a nonlinear least-squares regression algorithm is used of the form:

$$\min_{p \in R^n} g(x,p), l < p < u \tag{1}$$

© 2016 The Author(s)
Eurographics Proceedings © 2016 The Eurographics Association.

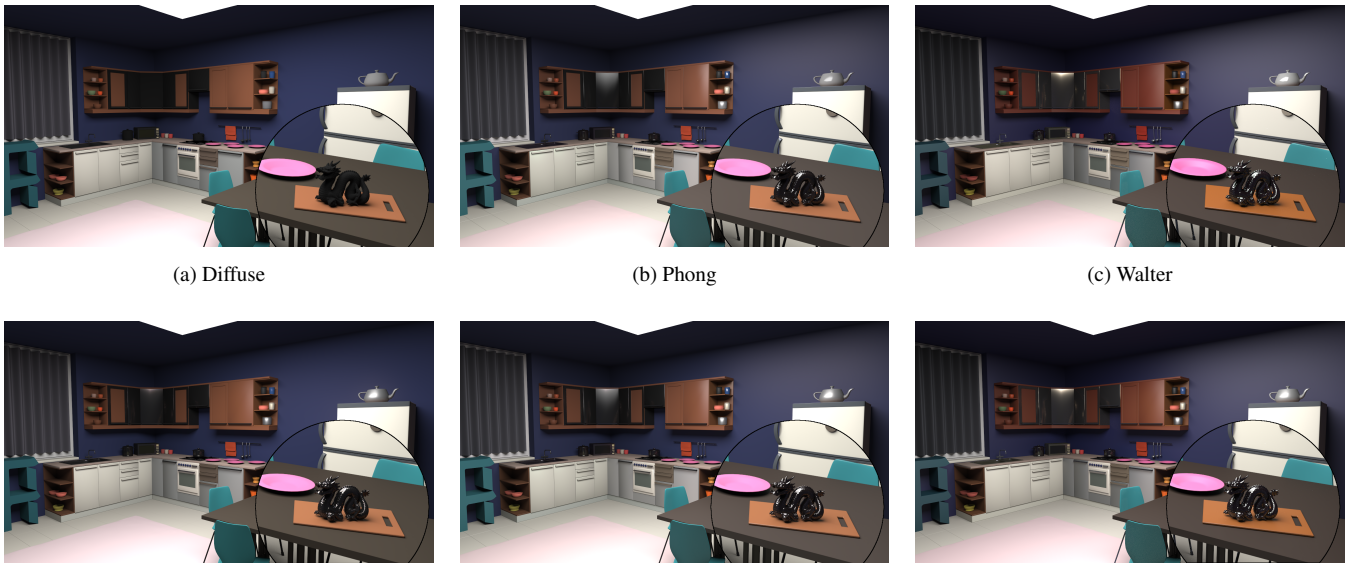

(a) Diffuse (b) Phong (c) Walter

(d) Sal-Diffuse (e) Sal-Phong (f) Sal-Walter

Figure 1: The Kitchen scene, rendered with the methods used in the experiment. The zoomed area shows the dragon sculpture, represented with the chrome MERL material

where x represents the fixed inputs to the function and p, l and u are vectors of length m containing the parameters for the BRDF in question and the lower and upper bounds of those parameters, respectively. The cost function, g, is defined as:

$$g(\omega_i, p_k) = \sum_{i=1}^{n} E_i^2 \tag{2}$$

where E_i is the squared error at the i^{th}data point on the k^{th} iteration and ω_i is the vector representing the direction of incoming light. To calculate the error a variant of the formulae described in the paper of Löw *et al.* is used.

$$E = \sin\theta_o \sqrt{y_m^2 - y_a^2} \tag{3}$$

where θ_o is the elevation angle of the view vector ω_o, expressed in spherical coordinates and y_m and y_a are the weighted outputs of the material and the analytical BRDF being fitted, respectively.

$$y_m = \ln(1 + \cos\theta_i f_r(\omega_i, \omega_o)) \tag{4}$$

$$y_a = \ln(1 + \cos\theta_i f_r(\omega_i, \omega_o, p)) \tag{5}$$

where θ_i is the elevation angle of the lighting vector ω_i and p is a vector containing the parameters for the analytical BRDF to be fitted.

In this weighting function the MERL material and analytical BRDF are queried and their returned RGB values are multiplied by the cosine of the incident vector's elevation, to reduce the effect of the poor data near grazing angles [NDM05]. Then this result is logged, in order to put an even weight on specular and non specular regions. The squared difference of each colour channel is then calculated.

In order to gather a manageable set of vectors to sample the material circles of increasing radii are projected from the unit disk onto the hemisphere centred around the perfect specular direction. Half of the projected circles have a small radius (approximately 0.2 on the unit disk), in order to capture the specular colour and glossiness and half have a larger radius, to capture the diffuse colour.

3.2.2. Scenes

The scenes chosen were all enclosed, indoor scenes. The three scenes were of a kitchen, a conference room and a lounge, which contained 68, 26 and 54 distinct materials, respectively. The reference images for each scene can be seen in Figure 2 and all images for the Kitchen scene can be seen in Figure 1 .

3.2.3. Salient Mixture Model

The mixture model stimuli are produced as hybrids of the MERL materials and analytical BRDFs. The saliency model employed is *Graph-Based Visual Saliency* (GBVS) [HKP06]. However this method is agnostic of the method and any other image-based visual saliency model could be used. An example saliency map, of the Kitchen scene, is shown in Figure 3. The mean saliency for the Conference, Kitchen and Lounge scenes was 0.1918, 0.2097 and 0.2165, respectively, where a value of one would represent a pixel with 100% salieny and a value of zero would represent a 0% saliency. This means that on average the reference method was sampled 19.18%, 20.97% and 21.65% of the time for the Conference, Kitchen and Lounge scenes respectively.

The salient mixture images are created by using the saliency of the given pixel to weight the BRDF that is sampled. Every time a

© 2016 The Author(s)
Eurographics Proceedings © 2016 The Eurographics Association.

(a) Conference (b) Kitchen (c) Lounge

Figure 2: Reference images, materials rendered using samples from the MERL database

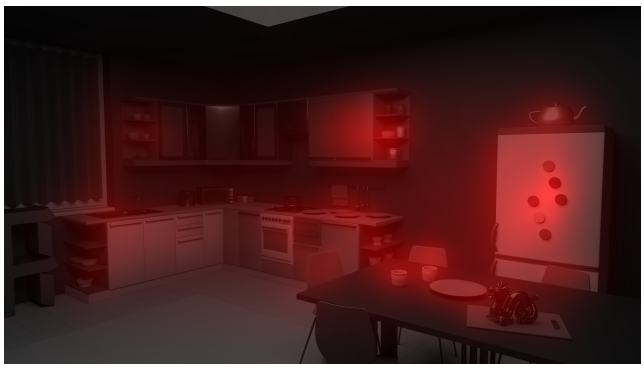

Figure 3: The Kitchen scene with the saliency of each pixel overlayed in red. Brighter red values represent a higher saliency value at that pixel.

surface is to be shaded, either the data-driven BRDF or an analytic model is selected based on sampling the normalised saliency value:

$$fr = \begin{cases} fr_{DATA} & \text{if } \xi < S(i) \\ fr_{ANALYTIC} & \text{otherwise} \end{cases} \quad (6)$$

where fr is the selected BRDF, fr_{DATA} is the relevent data driven BRDF, $fr_{ANALYTIC}$ is the analytic BRDF (Diffuse, Phong or Walter), $S(i)$ is the saliency normalized to [0,1] at pixel i, where the path originated, and ξ is a uniformly distributed random number between 0 and 1.

3.3. Participants

There were 30 total participants in this experiment, 4 female and 26 male from a variety of academic backgrounds. Only one possessed expert-level graphics knowledge. All participants had normal or corrected to normal vision.

3.4. Procedure

The experiment took take place in a room with low ambient light varying between 8.1 and 35.3 lux and the images viewed on a 55 inch HD monitor with the brightness and contrast set low to avoid eye fatigue, as per the ITU-R recommendation [Ass03]. The viewing distance was 2.1 metres to avoid discomfort, as per the same recommendations.

A sequential arrangement for the images was followed with the ground truth image displayed first followed by the stimulus. Both were displayed for a duration of three seconds [MTM12]. All stimuli were shown to all the participants.

This was followed by a five second break for the participants to rate the image before the next image in the series was displayed. The order of the images was randomised. Each experiment took around ten minutes to complete.

4. Results

In this section the results of the perceptual experiment are presented, analysed and discussed; as timings for rendering both the hybrid images and the images rendered with each method.

4.1. Timing

In order to assess the efficiency of each method, the time trials were run on a single CPU core. Timing results were conducted on a single thread using an Intel Xeon E5-2620 at 2.00GHz with 16gb of RAM on the Windows 7 64-bit operating system. The images were rendered at a resolution of 1920 × 1080 with one direct and one indirect lighting bounce. The timings for computing one sample per pixel are shown in Table 1. The images were rendered for 100 samples per pixel and the timings were then averaged to give the results below. The ratio of computation time vs. the reference is shown in the lower half of Table 1. This illustrates that computational gains can be expected using the saliency weighted methods.

4.2. Perceptual Tests

In order to analyse the similarity of images to the hidden reference the ratings given by participants are converted into distances between their rating of each condition and their rating of the relevant reference image, as recommended by Mantiuk *et al.* [MTM12].

$$d_{i,j,k} = r_{i,ref(k),k} - r_{i,j,k} \quad (7)$$

where d is a distance score, indicating the distance between a given participant's raw rating, r, of an image and their rating of the respective reference. Here i, j and k represent the observer, image and scene respectively.

© 2016 The Author(s)
Eurographics Proceedings © 2016 The Eurographics Association.

Table 1: Time taken (s) for one sample per pixel in a 1920×1080 image for each BRDF in each scene. Below are speedups of time taken to render with the reference method compared to the relevant method.

Scene	D	P	W	SD	SP	SW	R
Conference	16.27	16.95	18.08	16.69	17.20	17.71	18.63
Kitchen	13.11	13.96	15.00	13.8	14.20	14.69	15.62
Lounge	21.52	21.87	22.72	21.74	22.30	22.71	23.41
Conference	1.146	1.099	1.031	1.116	1.083	1.052	1.000
Kitchen	1.191	1.119	1.041	1.131	1.100	1.063	1.000
Lounge	1.088	1.070	1.030	1.076	1.050	1.030	1.000

Table 2: Mean Z-scores for each scene.

Scene	D	P	W	SD	SP	SW	R
Conference	0.670	0.404	0.281	0.009	0.319	0.270	-0.550
Kitchen	1.121	-0.127	-0.012	-0.328	-0.288	-0.440	-0.550
Lounge	1.039	-0.047	0.231	-0.590	-0.501	-0.364	-0.550
Overall	0.943	0.077	0.167	-0.303	-0.157	-0.178	-0.550

The scale used in this study, 1 to 100, is broad and different participants have different standards for what a low similarity score is. Therefore in order to more effectively compare results from different participants, Z-scores are calculated, see equation 8. This gives a measure of the extent to which participants thought a given image differed from the relevant reference image and sets the mean and standard deviation of each participant's ratings across all scenes to 0 and 1, respectively.

$$z_{i,j,k} = \frac{d_{i,j,k} - (\bar{d}_i)}{\sigma_i} \tag{8}$$

where σ_i is the standard deviation of the i^{th} observer's ratings across all images.

Table 3: Contrast comparisons between BRDFs in each scene. Coloured groupings indicate no significant differences across methods. *significant at $p < 0.01$

Scene	**BRDF**						**Kendall(W)**
Kitchen	*D*	*W*	*P*	*SP*	*SD*	*SW*	*R*	0.227*
Lounge	*D*	*W*	*P*	*SW*	*SP*	*SD*	*R*	0.267*
Conference	*D*	*P*	*W*	*SP*	*SW*	*SD*	*R*	0.213*
All	*D*	*W*	*P*	*SP*	*SW*	*SD*	*R*	0.397*

Figure 4 displays the mean scores for each method across all scenes. An average negative z-score reveals that a method was deemed to be above average by the participants, where zero is the average rating. The scores for the hidden reference are also included as participants could have rated the other images as more similar to the reference than itself. The reference images have scores of $\frac{-\bar{d}_i}{\sigma}$ for the i^{th} participant, by definition.

Results were analysed via two-way repeated measures analysis of variance (ANOVA) in a 7 (method) × 3 (scenes) factorial design. The main effect of scenes did not violate the assumption of

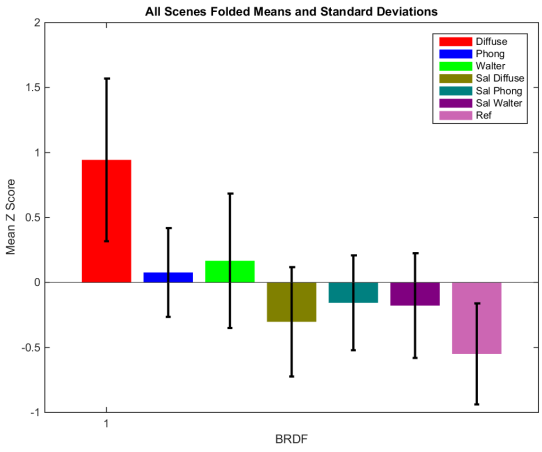

Figure 4: Aggregate scores and standard deviation for each BRDF across all participants and scenes

sphericity (Mauchly's Test of Sphericity, $p > 0.05$). The main effect of the scene did not produce significant differences, F(2, 384) = 3.018, $p > 0.05$, indicating no significant differences were found between scenes.

The main effect of method did not violate the assumption of sphericity (Mauchly's Test of Sphericity, P > 0.05) and was significant F(6, 384) = 29.445, $p < 0.01$. This suggests that the method used had a significant effect on the participants' ratings.

Kendall's co-efficient of agreement [KS39] was computed on the three scenes and on the collapsed overall scores to identify agreement across participants. Kendall's co-efficient gives a value of 0 when participants are in complete disagreement and 1 when in complete agreement. The results are shown in Table 3. The results are all considered significant ($p < 0.01$) indicating relative agreement in judging by the participants.

Pairwise comparisons with *Bonferroni corrections* were conducted to identify significant differences amongst the individual methods.

The results of these tests can be seen in Table 3. Coloured groupings represent methods between which a significant difference was not found. Results demonstrate groupings of the saliency methods and the other methods. As expected diffuse by itself performs poorer than all other methods, the Phong and Walter methods are grouped together as are all the mixed model methods. These groupings represent the lack of a discovered significant difference between how similar images created using different methods were to the reference image. Finally the reference scored better than the rest of the methods, though this difference was not found to be significant in some of the scenes, as seen in Table 3.

5. Discussion

The aggregate scores of all participants for each method are displayed in Figure 4. Here large negative results indicate a higher rating was given to the method. The pink bars on the right hand side of

© 2016 The Author(s)
Eurographics Proceedings © 2016 The Eurographics Association.

the figures represent the average ratings of the reference methods, for the relevant scene. The salient methods are significantly better than their counterparts, overall, but there was no significant difference to discriminate amongst them with the current set of captured data.

An examination of Table 2 shows the results for the Conference scene differ heavily from the Kitchen and Lounge scenes. In the conference scene all methods were rated lower than the reference and pairwise comparisons found a significant difference in the similarity ratings of the reference image and the hybrid images and between the hybrid images and their analytical counterparts, as can be seen in Table 3. In this experiment the saliency based methods did not exhibit significant difference in every scene and the overall average, suggesting they are of a similar quality. In addition Phong and Walter exhibited no significant difference while Diffuse has no correlation with any other method except in the conference scene.

The timings in Table 1 show that a decrease in computation cost can be expected across all methods. As can be seen in Table 2 the salient diffuse method achieved the greatest speed ups of the hybrid methods with savings of 7.6%, 11.6% and 13.1% for an average speed up of 10.8%. The Phong and Walter methods also outperformed the reference with average speed ups of 7.8% and 4.8% respectively. As Table 3 shows, visual perception of the scene is not significantly affected.

6. Conclusion and Future Work

This work has investigated whether analytical BRDFs can be substituted for data-driven models in a rendering pipeline without a noticeable effect to an observer. Three scenes and seven methods for representing materials were compared through the utilisation of a subjective study and statistical analysis, in two of the three scenes examined, significant differences were not found between the reference image and in all scenes the hybrids outperformed their non-salient counterparts. All three salient methods provided an improvement in computational performance with salient diffuse providing the largest mean computational saving (10.8%). This provides an indication that visual saliency can be used to improve computational performance by replacing expensive materials with cheaper analytic models in less visually important regions of the image.

This work could be expanded to approximate high quality BRDFs that incorporate physical approximations of diffraction [LKYU12] [HP15] and the polarisation of incident light [HTSG91]. These functions are more mathematically complicated and in some cases [HP15] require precomputation of the geometric shadowing and masking functions.

Acknowledgements

Tim Bradley is funded by the EPSRC. Chalmers & Debattista are partially supported by Royal Society Industrial Fellowships.

References

[AS00] ASHIKHMIN M., SHIRLEY P.: An anisotropic phong brdf model. *Journal of graphics tools 5*, 2 (2000), 25–32. 3

[Ass03] ASSEMBLY I. R.: *Methodology for the subjective assessment of the quality of television pictures*. International Telecommunication Union, 2003. 5

[Bli77] BLINN J. F.: Models of light reflection for computer synthesized pictures. In *ACM SIGGRAPH Computer Graphics* (1977), vol. 11, ACM, pp. 192–198. 3

[BLPW14] BRADY A., LAWRENCE J., PEERS P., WEIMER W.: genbrdf: Discovering new analytic brdfs with genetic programming, 2014. 2, 3

[BSH12] BAGHER M. M., SOLER C., HOLZSCHUCH N.: Accurate fitting of measured reflectances using a shifted gamma micro-facet distribution. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 1509–1518. 2, 3

[CCL02] CATER K., CHALMERS A., LEDDA P.: Selective quality rendering by exploiting human inattentional blindness: looking but not seeing. In *Proceedings of the ACM symposium on Virtual reality software and technology* (2002), ACM, pp. 17–24. 2

[CCW03] CATER K., CHALMERS A., WARD G.: Detail to attention: exploiting visual tasks for selective rendering. In *ACM International Conference Proceeding Series* (2003), vol. 44, pp. 270–280. 1, 2

[CDMPdS07] CHALMERS A., DEBATTISTA K., MASTOROPOULOU G., PAULO DOS SANTOS L.: There-reality: selective rendering in high fidelity virtual environments. *The International Journal of Virtual Reality 6*, 1 (2007), 1–10. 1, 2

[CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Transactions on Graphics (TOG) 1*, 1 (1982), 7–24. 2, 3

[DVGNK97] DANA K., VAN GINNEKEN B., NAYAR S., KOENDERINK J.: Columbia-utrecht reflectance and texture database, 1997. 2

[ENSB13] EISENACHER C., NICHOLS G., SELLE A., BURLEY B.: Sorted deferred shading for production path tracing. In *Computer Graphics Forum* (2013), vol. 32, Wiley Online Library, pp. 125–132. 1

[GDS14] GALEA S., DEBATTISTA K., SPINA S.: Gpu-based selective sparse sampling for interactive high-fidelity rendering. In *Games and Virtual Worlds for Serious Applications (VS-GAMES), 2014 6th International Conference on* (2014), IEEE, pp. 1–8. 1

[HE12] HEALEY C. G., ENNS J. T.: Attention and visual memory in visualization and computer graphics. *Visualization and Computer Graphics, IEEE Transactions on 18*, 7 (2012), 1170–1188. 2

[HKP06] HAREL J., KOCH C., PERONA P.: Graph-based visual saliency. In *Advances in neural information processing systems* (2006), pp. 545–552. 2, 4

[HP15] HOLZSCHUCH N., PACANOWSKI R.: *A physically accurate reflectance model combining reflection and diffraction*. PhD thesis, INRIA, 2015. 2, 3, 7

[HTSG91] HE X. D., TORRANCE K. E., SILLION F. X., GREENBERG D. P.: A comprehensive physical model for light reflection. In *ACM SIGGRAPH computer graphics* (1991), vol. 25, ACM, pp. 175–186. 3, 7

[IKN98] ITTI L., KOCH C., NIEBUR E.: A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 11 (1998), 1254–1259. 2

[KDCM14] KOULIERIS G. A., DRETTAKIS G., CUNNINGHAM D., MANIA K.: C-lod: Context-aware material level-of-detail applied to mobile graphics. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 41–49. 1, 2

[KS39] KENDALL M. G., SMITH B. B.: The problem of m rankings. *The annals of mathematical statistics 10*, 3 (1939), 275–287. 6

[KU87] KOCH C., ULLMAN S.: Shifts in selective visual attention: towards the underlying neural circuitry. In *Matters of intelligence*. Springer, 1987, pp. 115–141. 2

[KVDS96] KOENDERINK J. J., VAN DOORN A. J., STAVRIDI M.: Bidirectional reflection distribution function expressed in terms of surface

© 2016 The Author(s)
Eurographics Proceedings © 2016 The Eurographics Association.

scattering modes. In *European Conference on Computer Vision* (1996), Springer, pp. 28–39. 2

[LDC06] LONGHURST P., DEBATTISTA K., CHALMERS A.: A gpu based saliency map for high-fidelity selective rendering. In *Proceedings of the 4th international conference on Computer graphics, virtual reality, visualisation and interaction in Africa* (2006), ACM, pp. 21–29. 2

[LFTG97] LAFORTUNE E. P., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), ACM Press/Addison-Wesley Publishing Co., pp. 117–126. 1, 2, 3

[LKYU12] LÖW J., KRONANDER J., YNNERMAN A., UNGER J.: Brdf models for accurate and efficient rendering of glossy surfaces. *ACM Transactions on Graphics (TOG) 31*, 1 (2012), 9. 2, 3, 7

[Mca02] MCALLISTER D. K.: *A generalized surface appearance representation for computer graphics*. PhD thesis, University of North Carolina at Chapel Hill, 2002. 1

[MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics 22*, 3 (July 2003), 759–769. 1, 2

[MTM12] MANTIUK R. K., TOMASZEWSKA A., MANTIUK R.: Comparison of four subjective methods for image quality assessment. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 2478–2491. 5

[NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of brdf models. *Rendering Techniques 2005*, 16th (2005), 2. 2, 4

[ON95] OREN M., NAYAR S. K.: Generalization of the lambertian model and implications for machine vision. *International Journal of Computer Vision 14*, 3 (1995), 227–251. 2

[Pho75] PHONG B. T.: Illumination for computer generated pictures. *Communications of the ACM 18*, 6 (1975), 311–317. 1, 2, 3

[SCCD04] SUNDSTEDT V., CHALMERS A., CATER K., DEBATTISTA K.: Top-down visual attention for efficient rendering of task related scenes. In *VMV* (2004), pp. 209–216. 2

[SDL*05] SUNDSTEDT V., DEBATTISTA K., LONGHURST P., CHALMERS A., TROSCIANKO T.: Visual attention for efficient high-fidelity graphics. In *Proceedings of the 21st spring conference on Computer graphics* (2005), ACM, pp. 169–175. 2

[Smi67] SMITH B.: Geometrical shadowing of a random rough surface. *IEEE transactions on antennas and propagation 15*, 5 (1967), 668–671. 2

[WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques* (2007), Eurographics Association, pp. 195–206. 1, 2, 3

[WW07] WEIDLICH A., WILKIE A.: Arbitrarily layered micro-facet surfaces. In *Proceedings of the 5th international conference on Computer graphics and interactive techniques in Australia and Southeast Asia* (2007), ACM, pp. 171–178. 1

© 2016 The Author(s)
Eurographics Proceedings © 2016 The Eurographics Association.