

Selective BRDFs for High Fidelity Rendering

1025

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 are identified as possessing 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, the hybrid images were shown to have no observable significant difference from the 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 improvements in light transport algorithms, coherence for tracing and shading, filtering, and to reduce 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 3.30GB. This requires significant memory bandwidth during the frequent, and often incoherent, 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 results for these

more expensive models are likely to give further improvements.

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 these to several virtual scenes. A saliency map was generated for each scene and used to evaluate whether 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, the second section discusses the relevant work that is related to this project. The third section outlines the methodology that was used to prepare the scenes and reflectance models used in the experiment as well as the procedure and design of the perceptual experiment. The results of the perceptual tests are discussed in section five and section six concludes the paper.

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 [HE12].

The saliency map was proposed by Koch and Ullman [KU87]. Saliency maps are grayscale images, where the 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 smooth 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 eyetrackers or other methods available to validate their model by testing it with human subjects.

Cater *et al.* [CCL02] examined whether *inattentive 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 in effect 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 select regions of an image 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. Longhurst *et al.* introduced a system which did live anti-aliasing within a selective rendering framework, it functioned on the GPU.

Chalmers *et al.* [CDMPdS07] 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 to reduce the level of detail of subsurface light transport, with effects varying from subsurface scattering and refraction to a simple diffuse model (D), 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.

Two of the BRDFs used in this study are the Phong (P) BRDF [Pho75] and the Walter (W) 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 refraction through rough surfaces as well as reflection 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.

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 re-

placement for acquiring first-hand data [BSH12] [LKYU12] [BLPW14] [HP15].

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.

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

Another early focus was put on reducing the dimensionality of the BRDF space, to enable representing reflectance functions with fewer parameters. This easier, more intuitive tuning of parameters to match material appearance and smaller datasets to reduce the cost of measurement and storage of data-driven reflectance models. The most noteworthy of which is the Rusinkiewicz's half angle parametrisation [Rus98], which reduces the dimensionality of isotropic BRDFs, such as the MERL database from four to three. originally described by Neumann *et al.* [NNSK99].

Ngan *et al.* [NDM05] used Sequential Quadratic Programming to minimise a squared difference error metric. They evaluated the ability of seven analytical BRDFs to represent the MERL database. The Cook-Torrance, Ashikhmin-Shirley and He *et al.* fit the data well with minimal errors, simpler Blinn-Phong, Lafortune and Duer models performed poorly. The grazing angles were not included in the fitting, as the data was extrapolated when the MERL materials were captured. They established that BRDFs which used the half-angle parametrisation performed significantly better than those that used the difference angle parameterisation. 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. It 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 pre-computed 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] use genetic algorithms to generate new BRDFs from the MERL database, their system is called genBRDF. 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. It is Löw's fitting function that is adapted in the next section to fit to the methods that are used in the perceptual experiment.

3. Methodology

This work is motivated by the need to provide better overall performance for scenes that may not require detailed representations of the entire scene 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 give higher quality to 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. 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 permits comparisons with the other stimuli to identify perceivable differences across them.

The independent variable is 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 (The analytical BRDFs in question are a diffuse model, the Phong BRDF [Pho75], the Walter BRDF [WMLT07]), mixed saliency models consisting of the three chosen BRDFs mixed with the data-driven BRDF using a saliency map and a hidden reference.

The scenes chosen were all enclosed, indoor scenes. The scenes were enclosed to ensure that the number of indirect lighting bounces remained constant, for timing purposes. The three scenes were of a kitchen a conference room and a lounge.

The camera angle, BRDF fitting procedure, sampling algorithm, saliency model, image resolution, number of samples per pixel and viewing time for each image are constant across all scenes and BRDF mixtures. 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. 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.

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:

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

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:

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

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.

$$E = \sin \theta_o \sqrt{y_m^2 - y_a^2} \quad (3)$$

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

$$y_m = \ln(1 + \cos \theta_i f(\omega_i, \omega_o)) \quad (4)$$

$$y_a = \ln(1 + \cos \theta_i f(\omega_i, \omega_o, p)) \quad (5)$$

In this weighting function 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 over circles of increasing radii are projected from the unit disk onto the hemisphere centred around the perfect specular direction with half the circles projected having a small radius (approximately 0.2 on the unit disk), in order to capture the specular colour and glossiness and half having a larger radius, to capture the diffuse colour.

3.2.2. Scenes

The scenes chosen were all enclosed, indoor scenes. The scenes were enclosed to ensure that the number of indirect lighting bounces remained constant, to ensure consistent indirect lighting and ensure comparability of render times across scenes. The three scenes were of a kitchen a conference room and a lounge, containing 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 saliency model could be used. An example saliency map, of the conference scene, is shown in Figure ??.

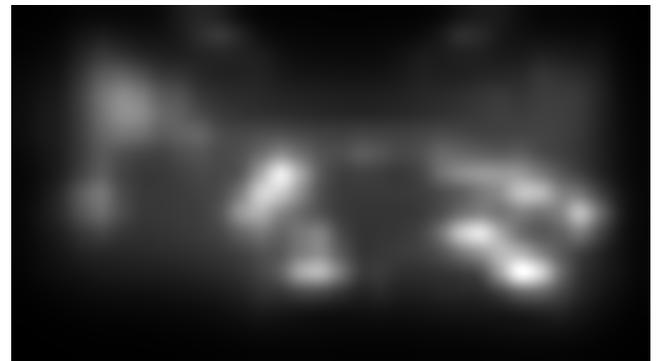


Figure 3: The saliency map generated with GBVS for the Conference scene.

The salient mixture images are created by using the saliency of the given pixel to weight the BRDF that is sampled. Every time a surface is to be shaded, either the data-driven BRDF or an analytic model is selected based on sampling the normalized saliency value:

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

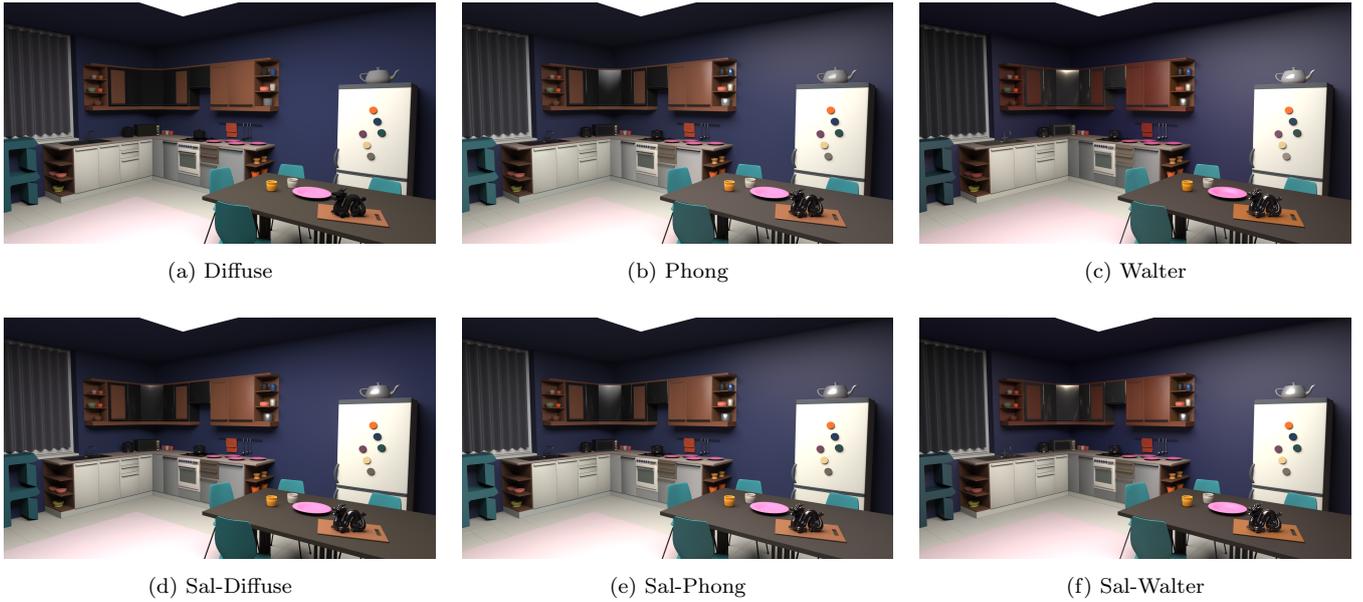


Figure 1: The Kitchen scene, rendered with the methods used in the experiment.



Figure 2: Reference images, rendered with the MERL database

where fr is the selected BRDF, fr_{DATA} is the relevant data driven BRDF, $fr_{ANALYTIC}$ is the analytic BRDF (diffuse, Phong or Walter), $S(i)$ is the normalized saliency 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 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 5 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 10 minutes to complete.

4. Results

In this section the results of the perceptual experiment are presented, analysed and discussed as well 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, to negate blocking is-

Table 1: Time taken (s) for one sample per pixel in a 1920×1080 image for each BRDF in each scene

Scene	D	P	W	SD	SP	SW	R
Conference	31.21	31.48	32.31	31.72	31.66	32.28	33.94
Kitchen	29.52	25.91	26.37	30.95	32.07	33.32	28.21
Lounge	42.48	42.02	42.57	42.51	42.84	42.86	42.95

Table 2: Ratios of time taken to render with the reference method in comparison to all other methods

Scene	D	P	W	SD	SP	SW
Conference	1.087	1.078	1.050	1.070	1.072	1.051
Kitchen	0.956	1.089	1.070	0.911	0.880	0.847
Lounge	1.011	1.022	1.009	1.010	1.003	1.002

sues and the same importance sampling method was used for creating each image. The images were rendered at a resolution of 1920 by 1080 with one direct and one indirect lighting bounce. The timings for computing one sample per pixel are shown in Table 1. The ratio of computation time vs the reference is shown in Table 2. This illustrates that on average, computational gains can be expected using the saliency weighted methods.

The recorded timings for the

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.

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. Doing so sets the mean and standard deviation of each participant’s ratings across all scenes to 0 and 1, respectively.

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

Figure 7 displays the mean scores for each BRDF across all scenes, the lower the score the more similar the set of images rendered with that particular BRDF are to the reference. The scores for the hidden reference are also included as theoretically participants could rate the other images as

Table 3: Contrast comparisons between BRDFs in each scene.

Scene	BRDF	pvalue	Kendall(W)
Kitchen	D (W (P (SP SD) SW) R)	< 0.05	0.227
Lounge	D (W (P (SW SP) SD) R)	< 0.05	0.267
Conference	D (P (W (SP SW) SD) R)	< 0.05	0.213
All	D (W (P (SP SW) SD) R)	< 0.05	0.397

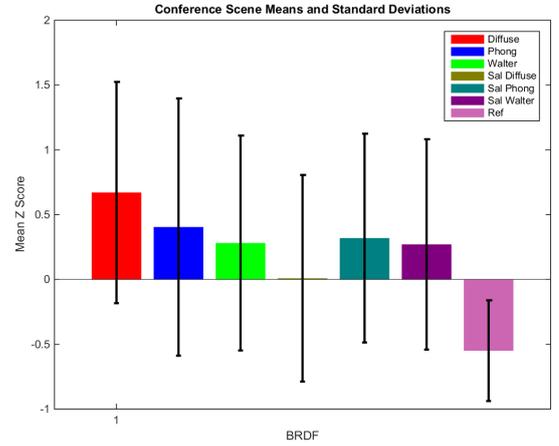


Figure 4: Mean Z scores and standard deviation for each BRDF across all participants for the Conference scene

more similar to the reference than itself, which will always have a score of $\frac{-d_i}{\sigma}$ for the i^{th} participant, by definition.

Results were analysed via repeated measures ANOVA in a 7 (method) × 3 (scenes) factorial design. The main effect of scenes did not violate the assumption of sphericity (Maulchy’s Test of Sphericity, $P > 0.05$) and did not pro-

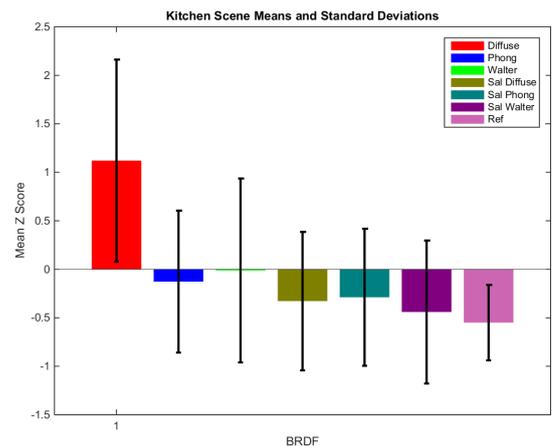


Figure 5: Mean Z scores and standard deviation for each BRDF across all participants for the Kitchen scene

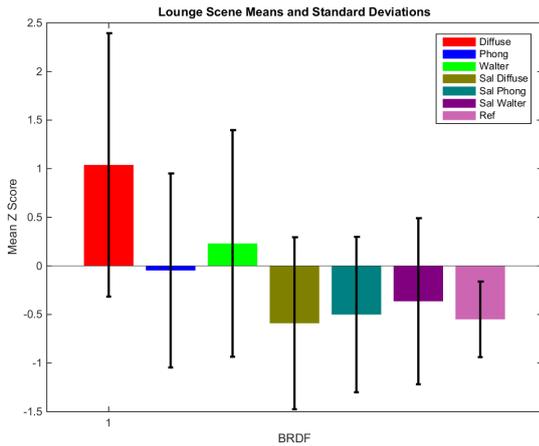


Figure 6: Mean Z scores and standard deviation for each BRDF across all participants for the Lounge scene

duce significant differences, $F(2, 384) > 0.05$, indicating no effect of scene upon the results. The main effect of method did not violate the assumption of sphericity (Maulchy’s Test of Sphericity, $P > 0.05$) and was significant $F(6, 384) = 29.445$, $P < 0.01$. Kendall’s co-efficient of Agreement was computed on the three scenes and on the folded 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 using 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 demonstrate lack of a significant difference amongst the methods. 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. Finally the reference is significantly better than the rest of the methods.

5. Discussion

As is visible in Figure 7, the Salient-Diffuse method was rated as the most similar to the reference and each hybrid method out-performed its counterpart. An examination of Figures 4, 5 and 6 show the results for the Conference scene differ heavily from the Kitchen and Lounge scenes. In the conference scene all methods were rated significantly 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. The saliency based methods exhibited no significant differ-

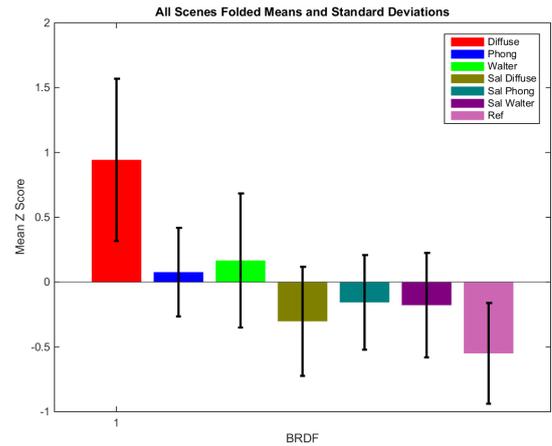


Figure 7: Mean Z scores and standard deviation for each BRDF across all participants and scenes

ence 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, while exhibiting the highest average rating for its hybrid method. This warrants further investigating as there may be a relationship between the simplicity of the model and its ability to create convincing hybrid images.

There is not enough data available to determine the cause of this and this will be the subject of future work. .

Timing results (see Tables 1 and 2) show that a decrease in computation cost can be expected, and as Table 3 shows, visual perception of the scene is not significantly affected.

6. Conclusion and Future Work

This work has investigated whether that 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, no significant difference was found between the reference image, which supports our hypothesis of The hybrid images and in all scenes the hybrids out performed their non-salient counterparts. 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.

References

- [Ass03] ASSEMBLY I. R.: *Methodology for the subjective assessment of the quality of television pictures*. International Telecommunication Union, 2003. 5
- [BLPW14] BRADY A., LAWRENCE J., PEERS P., WEIMER W.: genbrdf: Discovering new analytic brdfs with genetic programming, 2014. 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. 3
- [CCL02] CATER K., CHALMERS A., LEDDA P.: Selective quality rendering by exploiting human inattentive 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. 3
- [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. 4
- [HP15] HOLZSCHUCH N., PACANOWSKI R.: *A physically accurate reflectance model combining reflection and diffraction*. PhD thesis, INRIA, 2015. 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. 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
- [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., STAVRIDIS M.: Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *European Conference on Computer Vision* (1996), Springer, pp. 28–39. 3
- [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, 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. 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, 6
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of brdf models. *Rendering Techniques 2005*, 16th (2005), 2. 3, 4
- [NNSK99] NEUMANN L., NEUMANN A., SZIRMAY-KALOS L.: Reflectance models with fast importance sampling. In *Computer Graphics Forum* (1999), vol. 18, Wiley Online Library, pp. 249–265. 3
- [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. 3
- [Pho75] PHONG B. T.: Illumination for computer generated pictures. *Communications of the ACM* 18, 6 (1975), 311–317. 1, 2, 3
- [Rus98] RUSINKIEWICZ S. M.: A new change of variables for efficient brdf representation. In *Rendering techniques 98*. Springer, 1998, pp. 11–22. 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 microfacet 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