Pacific Graphics (2014)
J. Keyser, Y. J. Kim, and P. Wonka (Editors)

Short Papers

Perceptually Optimised Illumination for Seamless Composites

Andrew Chalmers[1], Jong Jin Choi[1], and Taehyun Rhee[1,2]

[1]Victoria University of Wellington
[2]taehyun.rhee@ecs.vuw.ac.nz

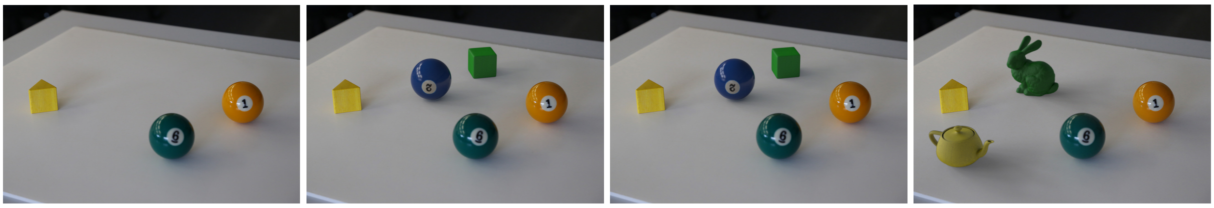

Figure 1: Composition with a photograph: the synthetic objects (blue ball, green block, bunny, and teapot) are illuminated by a radiance map (RM). From left: photograph, rendering with 5024x2512 HDR RM, and two examples of rendering with 320x160 LDR RM having expanded dynamic range using inverse tone mapping. The 2nd and 3rd images are visually equivalent, despite the large difference in RM resolution and dynamic range.

Abstract

Realistic illumination in composition is important for a seamless mixture between the virtual and real world objects in visual effects and mixed reality. The seamlessness is the measure of how perceivably apparent the synthetic object in the final composition is, and how indistinguishable it is from the photographed scene. Given that the ultimate receiver of image information is the human eye, the metric is determined by the Human Visual System (HVS). We conducted a series of psychophysical studies to observe and capture the thresholds of the HVS's ability to perceive illumination inconsistencies between the rendered subject and the photographed scene. Based on our observations, we find perceptually optimised thresholds for reducing resources across resolution and dynamic range of the radiance map (RM) for image based lighting (IBL). We evaluated our thresholds to illuminate virtual objects for seamless composition with photographed scenes.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. Introduction

Realistic mixture of synthetic objects with the photographed scene is one of the primary goals in mixed reality. In order to provide real world lighting information for synthetic objects using IBL, photographed RMs are commonly used. This requires the extra effort for capturing multi-exposure images to obtain the sufficient dynamic range. The process often requires extra time and effort or professional capturing devices. The final quality of the rendered output can vary greatly based on the parameter set considering the resolution and dynamic range of RMs, the complexity of geometry and materials, and parameters for rendering.

We can optimise these parameters for rendering while maintaining image quality defined by the HVS. Accordingly, recent computer graphics research has begun to focus on this effort [YCK*, KFB10]. The general framework of these studies is as follows: a reference image is rendered by a sophisticated rendering algorithm with unoptimised parameters. A test image is then rendered with perceptually optimised algorithms or parameters. The test image is then compared with the reference image to evaluate whether the optimisations have maintained image quality. The metric here is the ***visual equivalence*** (VE) between the images, where images have VE if they convey the same impression of the scene appearance, even if they are visibly different.

In the case of composition in visual effects and mixed re-

© The Eurographics Association 2014.

DOI: 10.2312/pgs.20141268

delivered by EUROGRAPHICS DIGITAL LIBRARY www.eg.org diglib.eg.org

ality, the reference image usually does not exist (e.g. synthesising aliens or spaceships in a photograph). To better replicate image composition in mixed reality, we develop a new framework for perceptual studies where there is no reference image. The aim of this study is different from previous perceptual studies because we measure ***seamlessness*** of the composition rather than the VE between two images. Using the self-referencing test, we conducted a series of psychophysical studies to capture the perceptual range of seamlessness in the HVS. Test images are generated by an IBL using high dynamic range (HDR) RMs that emulate the photographed real world scene. By manipulating the RM, we adjust lighting parameters such as light direction, intensity, and bit range in various environments. The changes are applied to render synthetic objects for composition. The composited output images are then shown to the participants, where they attempt to detect any abnormalities of illumination in a single composite image. Based on the analysis of the psychophysical experiments, we investigate which artefacts are salient to the HVS for perceiving seamlessness in composition.

Based on our observations from these tests, we explore the effects of optimising the HDR RMs in two aspects. First, we conducted a perceptual test using RMs of low dynamic range images (LDRIs), high dynamic range images (HDRIs), and a synthesised HDRI using inverse tone mapping operator (iTMO) from the LDRI. Our psychophysical study perceptually verified that the iTMO can reconstruct the dynamic range from the LDRI that is required for the IBL within the threshold that the HVS cannot perceive. Furthermore, we reduce the resolution of the RM based on the fact that the HVS is insensitive to changes in light direction and intensity. Our results show that the size of the RM can be reduced significantly while maintaining the seamlessness of the final composition. This reduction directly results in memory savings; with an average of 99% savings of a general studio setup. Optimising memory is one of the important issues in mobile graphics.

2. Related Work

Illumination composition: Compositing synthetic objects with a photograph has been an active research topic in computer graphics and augmented reality [NHIN, FGR]. Although previous works related to tracking and registration need to be fully considered, due to the large amount of literature, we will focus only to the works mainly contributing to illumination and rendering for realistic composition. In order to render 3D objects for image composition, we need to know the real-world light information. There are a variety of ways to capture real world lighting in an environment for rendering, all ranging in cost of setup time and complexity [MPK, LEN]. Composition often uses the method of HDRIs [DM] to capture the radiance of the environment, giving realistic lighting and shadows for the synthetic object. Advancements in this area can be found in [Gro, KTM*] for composition, and [SJW*] for HDR capture.

Perceptual studies in this domain is limited and has recently gained interest. Such focuses include illumination direction, object orientation, materials and textures, composition, and user studies [LMSSG, KHFH].

Visual perception for visual equivalence: Perceptually based rendering increases efficiency while maintaining perceptible image quality. This approach exploits the nature of the HVS [OHM*, BCFW]. Ramanarayanan et al. [RFWB] introduced the notion of VE, which has been considered in other works [FRWB, KFB10].

Perceptually based rendering and optimisation: Perceptually based rendering involves optimising a rendering system while satisfying the HVS. This avoids the one-to-one simulation of the physical world, by only making computations that are necessary for image detail that the HVS is able to perceive. Works include perceptuals studies to optmise rendering parameters for radiosity [YCK*] and shadows for interior lighting [NKO].

HVS for seamlessness: Research that conducts perceptual studies without explicit reference images is limited, with focuses on shape and material [VLD] and resolution [CCR]. Alternatively, we provide a study of varying scenes, object shapes, and materials designed for composition while evaluating the effects of converting LDR to HDR - showing that these different properties lead to different approximations.

HVS for inverse tone mapping: IBL rendering typically makes use of HDRIs for RMs. Converting LDR to HDR can be useful for saving time by not capturing multiple exposure levels, as well as converting legacy images to HDR. There are simple methods for expanding the range of LDR images [Lan], as well as more advanced techniques [BLD*]. The advanced techniques were evaluated using the HDR visible differences predictor (VDP) metric [VDP] instead of a psychophysical study for composition. We refer to [BDA*] for an extensive survey on LDR to HDR expansion.

3. HVS in Local Illumination Changes

Previous studies that observe the HVS in composited scenes are limited. A novel experiment setup for psychophysical tests is required to measure seamlessness in lighting composition. We make *local illumination changes* to composited objects, where the local illumination change refers to the altered illumination properties of the composited object (as opposed to the unaltered illumination of the other objects in the scene). In addition to the psychophysical test in changes of light direction and intensity [CCR], we ran several extra experiments measuring the visibility of local illumination changes, namely the dynamic range of RMs, and the influence of rendering with a HDR image which has been converted from a LDR image.

© The Eurographics Association 2014.

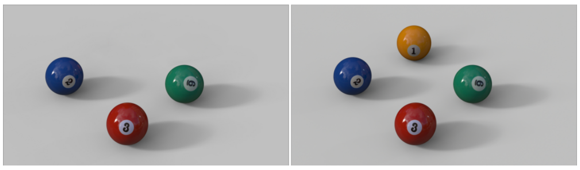

Figure 2: Examples of RM modifications in the user study, where the left is an IBL rendering with the original HDR RM – this acts as a simulation of a photograph. The right has a "synthetic" object composited into the photograph under improper illumination. The synthetic object is illuminated with an 8-bit RM.

3.1. Experiment Setup

The aim of our experiment is to observe human perception relating to local illumination changes without reference images. The images of the background scene are captured by a fish-eye-lens camera at multiple exposure levels and angles. Then, the images are stitched together to form a 360-degree 16-bit HDR RM with a resolution of 5024 x 2512 (a common resolution in film production). Using the HDR RM, the background scenes are rendered by IBL using pathtracing [PH, Faj]. The background scenes shown in Figure 2 (left column) mimics the photograph in our test. A synthetic object is lit under the modified RM to simulate improper lighting. Then, it is composited into the background scene using differential rendering [DM]. In a preliminary user test, we used a large image set and a small number of participants to narrow down the test set. After finding the threshold range to focus on, we increased the number of participants and decreased the number of test images to collect more data on a meaningful range. The rendered images are a resolution of 1280x720. An example of a composition with a modified parameter of an IBL can be seen in Figure 2 (right column).

3.2. Stimuli

In order to focus our test on the illumination changes, we limit the object geometry and materials. Each test image includes four objects having similar shape and topology. Objects include billiard balls and wooden blocks. Cubes were modelled as wooden blocks, with mostly diffuse material. Billiard balls were modelled as spheres with a glossy-specular material, since specular components are also an important source of information for visual perception [FTA]. Previous research has shown that spheres are the least discriminating shape with respect to material properties [VLD]. Our choice of cube is based on the fact that it contributes a more complex shadow than the sphere.

The tests are subdivided into different lighting environments using three HDR RMs representing real world environments:

- **Sunny RM:** Outdoor scene with dominant sun light.
- **Overcast RM:** Outdoor scene, under partial shelter during wet weather and an overcast sky.
- **Indoor RM:** Indoor scene with multiple light sources emitting from the windows and indoor artificial lights.

The RMs are shown in Figure 4. The tests are further split into categories based on the lighting parameters. The variation of the lighting parameters of the virtual object are simulated by the manipulation of the RMs:

- **8-Bit LDR:** Test objects are rendered with a LDR RM.
- **Illumination intensity:** Test objects are rendered with different values of intensity for the HDR RM, where intensity is defined as a scalar of the light intensity from the RM.
- **Illumination direction:** The test objects are rendered with the lighting direction changed in increments of five degrees across the azimuth angle, from the 0 to 45 degrees. The range is filtered based on the pilot user tests.
- **LDR to HDR:** The test objects are rendered with an expanded LDR to HDR RM.
- **Reduced resolution:** The test objects are rendered with a smaller RM.

3.3. Procedure

We produce the set of IBL adjustments for each category, which is a combination of two shapes, two materials, and three lighting environments. See Table 1 for the number of variations that were tested.

Table 1: The set of stimuli in the perceptual study

	Sunny	Overcast	Indoor
Direction	10 diffuse 10 specular	10 diffuse 10 specular	10 diffuse 10 specular
Intensity	6 diffuse 6 specular	6 diffuse 6 specular	6 diffuse 6 specular
LDR	diffuse specular	diffuse specular	diffuse specular
LDR to HDR	diffuse specular	diffuse specular	diffuse specular
Resolution	7 diffuse 7 specular	7 diffuse 7 specular	7 diffuse 7 specular

The test images are shown to the participants without an explicit reference image. The order of images are randomised and include images with no adjustments. The randomness, range of scenes, and object types reduces the likelihood that participants can detect patterns. The participants are asked if they could identify the areas which have been adjusted within the test image, and how noticeable this adjustment is. The categories are labelled as follows: 0. (Not noticeable), 1. (Slightly noticeable), 2. (Moderately noticeable), 3. (Very noticeable), and 4. (Extremely noticeable), adapted from [MCC], and participants are asked which object has the modification. If the participants cannot correctly identify the adjustment in the correct part of the image, we

© The Eurographics Association 2014.

consider that they cannot distinguish the change in illumination. Our user study is made up of 20 participants, all of whom are shown all images. All participants have normal or corrected-to-normal vision.

3.4. HVS Observation and Analysis

In this section, we observed the basic characteristics of the HVS in image composition while varying light direction, intensity, and the dynamic range of the RM. The results motivate our next step to optimise the RM, as discussed in section 4 and 5.

Illumination direction: We tested the HVS's ability to perceive light direction changes. Light direction was investigated by rotating the RM from 0 to 45 degrees in 5-degree increments. The results are shown in Figure 3a, showing the range between 5 and 30 degrees. For the light direction test, it starts becoming noticeable from 15 degrees of divergence, it becomes increasingly noticeable from 30 degrees onward. These abnormalities are typically recognised by the angle of the specular highlight or shadow area.

Illumination intensity: We tested the ability of the HVS to perceive light intensity changes. It was simulated by reducing the overall intensity values of the RM in increments of half the intensity value. The results are shown in Figure 3b. Based on the experiment, participants could not observe any abnormality while decreasing intensity by 50% (we skip this range and display the rest in Figure 3b). After 50% reduction, intensity changes show a steady increase in noticeability and are more apparent than the shadow changes.

8-bit LDR: We perform additional tests to perceive local illumination changes while altering the bit range of the RM from HDR to LDR. The orignal 16-bit HDRIs are converted to 8-bit LDRIs (with consideration for gamma correction). The LDRIs are then used for RM to render test objects. Using the sunny and indoor RM, many participants preceived slight artefacts. The majority of people could not notice artefacts in local illumination changes rendered by the LDR overcast RM. In the overcast RM, the distribution of the intensity in the original HDRI is mostly below the 8-bit range. Therefore, in the case of overcast scenes such as cloudy and rainy days, it is perceptually acceptable to use LDR as the RM.

4. HVS for Inverse Tone-Mapped Radiance Map

We conduct psychophysical tests to evaluate an iTMO that expands the dynamic range of an LDR RM. In this study, we utilised iTMO of [Lan], where the expanded luminance L_o is defined as:

$$L_o(\mathbf{x}) = \begin{cases} L_\gamma(\mathbf{x})(1-i) + L_\gamma(\mathbf{x})\varepsilon i, & \text{if } L_\gamma(\mathbf{x}) \geq \upsilon \\ L_\gamma(\mathbf{x}), & \text{otherwise} \end{cases}$$

$$i = \left(\frac{L_\gamma(\mathbf{x}) - \upsilon}{1-\upsilon}\right)^\alpha$$

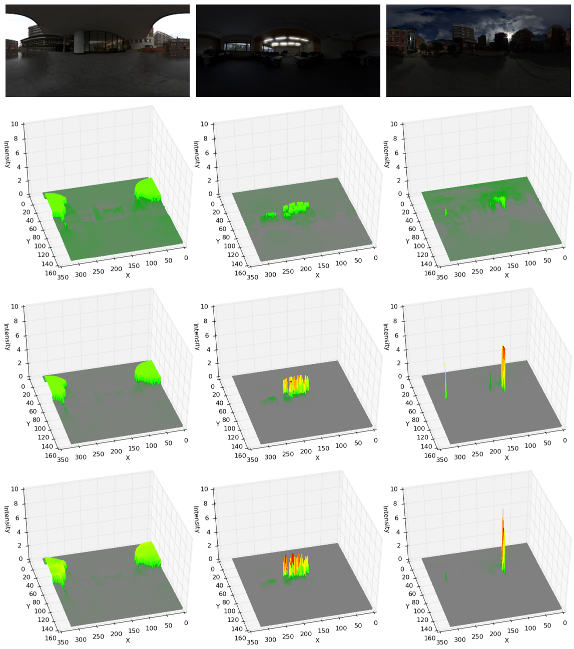

Figure 4: The dynamic range of each scene. Column 1, 2 and 3 is the overcast, indoor and sunny scene respectively. Row 1 is the radiance map, and rows 2, 3 and 4 is the dynamic range of the LDR, LDR to HDR, and HDR respectively. The LDR luminance values have been normalised between 0 and 1.

Where L_γ is the inverse gamma corrected function of the LDR gamma-encoded luminance value, ε is the expanded range, α is the exponent falloff, and υ is a threshold value defining the starting luminance for expansion. The value ε is approximated experimentally.

We render the background scene with the original HDR RM. The original HDR RM is then converted into a gamma encoded LDR RM, which is then converted back into a HDR RM using the iTMO. We render the synthetic objects using the iTMO converted RM. Then, they are composited into the background scene. All RMs are of high resolution (5024x2512).

Results of the test are shown in Figure 3d. We find that the participants in the survey were generally unable to distinguish which object had been illuminated with the alternate RM. This shows that the iTMO produced a sufficient composite that seamlessly blends with the rest of the scene.

We observe that the dynamic range in the overcast scene is quite limited, thus the LDR image is sufficient in capturing the lighting information. The indoor scene has multiple bright light sources, thus making it necessary to capture the illumination properties using a wider dynamic range. Further, we see differences in illumination from the various

© The Eurographics Association 2014.

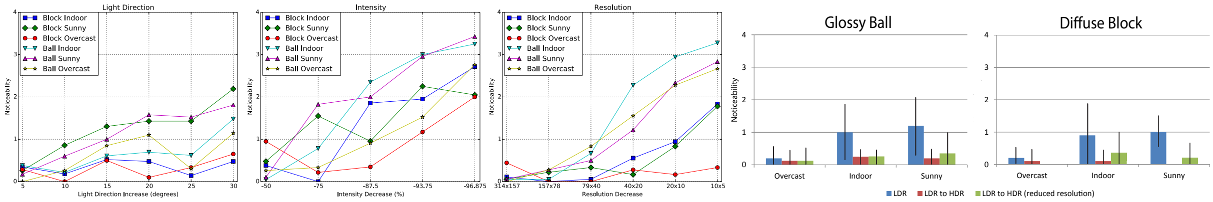

(a) (b) (c) (d)

Figure 3: Observation of the HVS. From left: illumination direction, intensity, resolution reduction, LDR to HDR comparison for glossy balls, and LDR to HDR comparison for diffuse blocks. The noticeability scale is the mean score of the participants. The graph is focused in the range of sensitivity, where illumination details were begining to be detected by the HVS.

light sources (the immediate light bulbs contrast with the dim light coming through the window). The iTMO matches similarly to the ground truth in this case. Finally, the outdoor scene has a clear sky with a distinct light source (the sun). The dynamic range is very wide, and the iTMO closely matches this range - though some of the detail is lost, where the peak is flattened compared to the HDR image which has a distinct peak. This is due to the fact that the LDR image has a wide area of the sun capped at 255, thus expanding this range maintains the flattened peak. The reflection in the window has the dynamic range boosted as well, causing a second peak (which does not match the ground truth). This is because of the same problem - the values in the reflection are capped at 255. Whereas the wider dynamic range can differentiate between bright values in the reflection of the window, and the brighter values in the sun. Though these artefacts are not perceptible in the final rendering for composition as shown by the results.

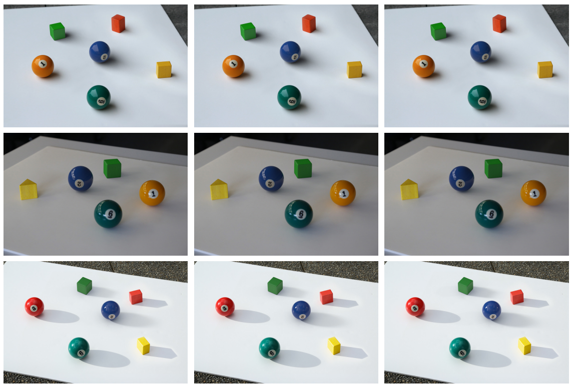

Figure 5: The final composition with a photograph, where each column from the left is illuminated with a HDR 5024x2512, HDR 320x160, and LDR to HDR 320x160 resolution RM. The blue ball and green block are synthetic.

5. Perceptually Optimised Radiance Maps

Based on our observation in section 3, we conduct more studies to find the perceptual threshold to optimise the resolution of the RMs. Since we found that there is some threshold that can be exploited with direction and intensity changes, reducing the RM has similar effects. For example, a small RM will create offset shadows or narrow shadows. We conduct the same experiment as outlined in section 3 to find optimal thresholds for the resolution of the RMs. We produce 7 variations of each category with a reduced RM resolution, and one extra image with no adjustments in each category. This is a total of 48 images to run the experiment with. We had a total of 30 participants in the survey for the test. Based on the results shown in Figure 3c, we chose the RM resolution that was optimised before being noticeable. The resolution 320x160 still maintains the same visual quality as a RM with a resolution of 5024x2512. This is a saving from approximately 43mb to 183kb. We also find that diffuse objects except the sunny day can be further reduced down to 80x40 resolution.

As shown in section 3, reducing the dynamic range of a RM from 16 to 8-bit causes significant loss in quality for composition. However, the iTMO can nicely reconstruct the dynamic range of LDR RMs as shown in section 4. Based on these results, we performed a final psychophysical test to find the optimal resolution range for RMs which are reconstruced by the iTMO. As shown in Figure 3d, most participants could not notice abnormalities in the final rendering and composition using the reconstructed LDR RMs by the iTMO for a RM resolution of 320x160. These results contribute by saving memory to store the RM, as well as the effort in capturing a LDR RM.

6. Composition with a Photograph

The results are tested with real photographs. We have applied our observations and thresholds to optimise the RMs. The RMs were then used to illuminate the synthetic objects, which are then composited with the real photographs as shown in Figure 5. The background scene consists of a mixture of wooden blocks and billiard balls. The photographs of the background scenes are taken in the same lighting conditions as the HDR RMs. As shown in the Figure 5, the optimised RMs provides sufficient lighting for the composition. The result is visually equivalent to the IBL rendering using the original HDR RM.

© The Eurographics Association 2014.

7. Conclusion

The aim of this paper is to observe the nature of the HVS to perceive seamlessness of image composites while performing a series of psychophysical tests without a reference image. The obervation lead to further experiments to quantify perceptually optimised thresholds for reducing resources for IBL with respect to both dynamic range and image resolution. Our results show that the iTMO can accurately reconstruct the dynamic range required for IBL in image composition in the range of diffuse to glossy-specular material types.

Our results are tested with real photographs and showed that the quality of the final IBL rendering using our optimised RM is visually equivalent to the IBL rendering using the original HDR RM. In the test, we explicitly used the RM as the only light source in the scene. Therefore, even in the IBL rendering with the original HDR RM, some artefacts (e.g. shadow edge and highlight spot) are slightly noticeable in the composition with strong sunlight. In film production, this is often solved by artists by adding artificial lights and tuning them until they achieve satisfactory results. Note that the number of test images in Section 3 is 155. This number should be limited to manage user test time in the pyschophysical study. Our perceptual test framework can be extented to identify optimal ranges of other parameters in future research.

References

[BCFW] BARTZ D., CUNNINGHAM D., FISCHER J., WALLRAVEN C.: The Role of Perception for Computer Graphics. Eurographics State-of-the-Art-Reports 2008, pp. 65–86. 2

[BDA*] BANTERLE F., DEBATTISTA K., ARTUSI A., PATTANAIK S., MYSZKOWSKI K., LEDDA P., CHALMERS A.: High dynamic range imaging and low dynamic range expansion for generating hdr content. Computer Graphics Forum, 28 2009. 2

[BLD*] BANTERLE F., LEDDA P., DEBATTISTA K., CHALMERS A., BLOJ M.: A framework for inverse tone mapping. *The Visual Computer 2007 23*, 7, 467–478. 2

[CCR] CHALMERS A., CHOI J. J., RHEE T.: Perceptually based radiance map for realistic composition. Proc. of Image and Vision Computing New Zealand 2013, pp. 172 – 177. 2

[DM] DEBEVEC P. E., MALIK J.: Recovering high dynamic range radiance maps from photographs. SIGGRAPH 1997, ACM, pp. 369–378. 2, 3

[Faj] FAJARDO M.: *Arnold Software 2013*. 3

[FGR] FOURNIER A., GUNAWAN A. S., ROMANZIN C.: *Common Illumination between Real and Computer Generated Scenes 1992*. Tech. rep. 2

[FRWB] FERWERDA J. A., RAMANARAYANAN G., WALTER B., BALA K.: Visual equivalence: an object-based approach to image quality. Proc. of IS&T 16th Color Imaging Conference. 2008, pp. 347–354. 2

[FTA] FLEMING R. W., TORRALBA A., ADELSON E. H.: Specular reflections and the perception of shape 2004. *J. Vis. 4*, 9, 798–820. 3

[Gro] GROSCH T.: Differential photon mapping: Consistent augmentation of photographs with correction of all light paths. Eurographics 2005. 2

[KFB10] KŘIVÁNEK J., FERWERDA J. A., BALA K.: Effects of global illumination approximations on material appearance. SIGGRAPH 2010, ACM, pp. 112:1–112:10. 1, 2

[KHFH] KARSCH K., HEDAU V., FORSYTH D., HOIEM D.: Rendering synthetic objects into legacy photographs. SIGGRAPH ASIA 2011, ACM, pp. 157:1–157:12. 2

[KTM*] KNECHT M., TRAXLER C., MATTAUSCH O., PURGATHOFER W., WIMMER M.: Differential instant radiosity for mixed reality. Symposium on Mixed and Augmented Reality, ISMAR 2010, pp. 99–107. 2

[Lan] LANDIS H.: Production-ready global illumination. SIGGRAPH Course Notes 16 2002. 2, 4

[LEN] LALONDE J.-F., EFROS A. A., NARASIMHAN S. G.: Webcam clip art: appearance and illuminant transfer from time-lapse sequences. SIGGRAPH ASIA 2009, ACM, pp. 131:1–131:10. 2

[LMSSG] LOPEZ-MORENO J., SUNDSTEDT V., SANGORRIN F., GUTIERREZ D.: Measuring the perception of light inconsistencies. Symposium on Applied Perception 2010, ACM, pp. 25–32. 2

[MCC] MEILGAARD M., CIVILLE G., CARR B.: *Sensory evaluation techniques 1991*. CRC Press. 3

[MPK] MURY A. A., PONT S. C., KOENDERINK J. J.: Representing the light field in finite three-dimensional spaces from sparse discrete samples. *Applied Optics 2009 48*, 3, 450–457. 2

[NHIN] NAKAMAE E., HARADA K., ISHIZAKI T., NISHITA T.: A montage method: the overlaying of the computer generated images onto a background photograph. SIGGRAPH 1986, ACM, pp. 207–214. 2

[NKO] NAKANO G., KITAHARA I., OHTA Y.: Generating perceptually-correct shadows for mixed reality. Symposium on Mixed and Augmented Reality, ISMAR 2008, pp. 173–174. 2

[OHM*] O'SULLIVAN C., HOWLETT S., MORVAN Y., MCDONNELL R., O'CONOR K.: Perceptually Adaptive Graphics. Eurographics State of the Art Reports 2004, pp. 141 – 164. 2

[PH] PHARR M., HUMPHREYS G.: *Physically Based Rendering 2010*, 2nd ed. Morgan Kaufmann Publishers Inc. 3

[RFWB] RAMANARAYANAN G., FERWERDA J., WALTER B., BALA K.: Visual equivalence: towards a new standard for image fidelity. SIGGRAPH 2007, ACM. 2

[SJW*] STUMPFEL J., JONES A., WENGER A., TCHOU C., HAWKINS T., DEBEVEC P.: Direct hdr capture of the sun and sky. SIGGRAPH 2006, ACM. 2

[VDP] VDP: http://www.mpi-sb.mpg.de/resources/hdr/vdp/. 2

[VLD] VANGORP P., LAURIJSSEN J., DUTRÉ P.: The influence of shape on the perception of material reflectance. SIGGRAPH 2007, ACM. 2, 3

[YCK*] YU I., COX A., KIM M. H., RITSCHEL T., GROSCH T., DACHSBACHER C., KAUTZ J.: Perceptual influence of approximate visibility in indirect illumination. *ACM Transactions on Applied Perception 2009 6*, 4, 24:1–24:14. 1, 2

© The Eurographics Association 2014.