

Validation of Color Managed 3D Appearance Acquisition

Michael Goesele, Hendrik P. A. Lensch, Hans-Peter Seidel
MPI Informatik, Saarbrücken, Germany

Abstract

Image-based appearance acquisition algorithms are able to generate realistic 3D models of real objects but have previously not taken care of calibrated color space. We integrate a color managed high-dynamic range imaging technique into a recent appearance acquisition algorithm and generate models in CIE XYZ color space. We compare the final models with spectrophotometric measurements and compute difference images between renderings and ground truth images. Displayed renderings and printouts are compared to the original objects under identical illumination conditions to evaluate and validate the complete appearance reproduction pipeline. Working in CIE XYZ color space allows for expressing the perceivable differences in a standardized measure.

1. Introduction

A visually realistic 3D model of a real object represents both its geometry as well as its surface properties. Traditionally, a diffuse texture is pasted onto the geometry model but realism is limited as a diffuse texture cannot represent an object's true reflection properties. Additionally, specular reflections make capturing a diffuse texture challenging and highlights caused by the illumination are often baked into the surface leading to visible artifacts. In contrast to diffuse textures, a bidirectional reflectance distribution function (BRDF) [16] describes how light is reflected off an object's surface. A BRDF can faithfully reproduce the appearance of an object's surface including local lighting and view-dependent effects for arbitrary viewing and lighting directions (see e.g. [5] for a more detailed description of the BRDF and its properties). Various physically based or empirical BRDF models such as the Cook-Torrance model [2], the He model [7], the Ward model [19], or the Lafortune model [13] have been proposed to represent the reflection properties of objects, mainly in the context of rendering. Several image-based acquisition techniques [15, 14] have been introduced in recent years to determine the parameters of these model for real objects.

Two acquired models with spatially varying BRDFs captured with and without applying color calibration are shown in Figure 1. The images illustrate the problem of



Figure 1: Calibrated (left) and uncalibrated (right) carafe dataset. Both datasets are generated from the same input images – once using the color managed work-flow and once using standard techniques. Variations in the highlights are caused by the optimization process used to generate the models.

correct color reproduction: The two images show renderings of BRDF models generated from the same input data that differ mainly in color. Given any of the two images (or even both), it is impossible to determine the color of the original object. The objective of our approach is not only to acquire a BRDF model with correct colors but also to ensure correct reproduction of the rendered BRDF model on arbitrary output devices, and to measure the differences in a standardized way.

1.1. Contribution

In this paper, we integrate ICC-based color management into an image-based BRDF acquisition system in order to acquire a high quality model. Our goal is in particular to achieve best possible color and appearance reproduction with arbitrary input and output devices using the color management mechanism (i.e., without manual tuning of system parameters). The input images are converted into high dynamic range (HDR) images using a color managed technique similar to Goesele et al. [6]. The whole BRDF computation is performed in CIE XYZ color space. We validate our approach by comparing the model at selected surface points with spectral measurements. We furthermore compute difference images (ΔE) between renderings of the model and real views thereby checking the integrity of both the BRDF modeling as well as the color management work-flow. Using color managed output devices, we are able to generate synthetic renderings of real objects that

match the original closely both in color as well as in specular. Furthermore, we can extract a “perfect diffuse texture” from the BRDF data by virtually evaluating it under standard 45/0 measurement conditions [8].

2. Previous Work

The appearance of a surface is traditionally measured using a range of specialized devices [9]. These devices can usually obtain only a single sample per measurement and are therefore not well suited to acquire the spatially varying appearance of a whole object. Image-based techniques overcome this limitation by simultaneously acquiring multiple samples with a single shot of a digital camera. Marschner et al. [15] used such an approach to acquire a single BRDF per object. Based on this technique, Lensch et al. [14] are able to capture a BRDF per surface point (spatially varying BRDF) yielding a detailed surface representation. We follow in this paper their approach and extend it with an ICC-based color management.

Capturing specular highlights on objects as well as the diffuse reflection require high-dynamic range (HDR) imaging techniques [3, 17]. These techniques generally recover a response curve of the camera system used and combine the resulting linearized input images with different exposure settings into a single HDR image. Besides a simple white balancing step, no color correction can be used. Goesele et al. [6] used a color management system to linearize the input images and generate HDR images in a defined color space. We use an improved version of this technique to acquire the HDR images.

3. Acquisition and BRDF Generation

We employ in this work the BRDF measurement technique by Lensch et al. [14] which we will review here only shortly. Please refer to the original publication for an in-depth description of the approach.

The target object is illuminated by a point light source in a dark and low reflecting room. A digital camera is used to acquire HDR images of the object from different view-points and for varying light source positions. Several calibration steps ensure that the position of both camera and light source relative to the object are known. Furthermore, a geometry model of the target object is generated using 3D scanning technology. Given this input data, the surface is clustered in different basis material clusters and the parameters for the analytical Lafortune BRDF model [12]

$$f(\vec{u}, \vec{v}) = \rho_d + \sum_i [C_{x,i}(u_x v_x + u_y v_y) + C_{z,i} u_z v_z]^{N_i}$$

are determined for each cluster. The actual measurements per surface point are then projected into a basis formed by these per-cluster BRDFs in order to recover a truly spatially varying and highly detailed model of the reflection properties.

3.1. BRDFs in CIE XYZ Color Space

The Lafortune BRDF model is originally defined in linear RGB color space which is closely related to the CIE XYZ color space used in the profile connection space (PCS) of an ICC-based color management system. The similarity of the color spaces makes BRDF generation in CIE XYZ straightforward. The different topologies of the color spaces (e.g. different definitions of the luminance channel) require however some parameter changes in the non-linear fitting of the BRDF model to the input data where at some stages chromaticity and luminance are considered separately.

4. Color Managed HDR Imaging

The acquisition was performed with a Kodak DCS 560 digital camera and a K5600 Joker Bug 800 W HMI lamp as point light source. To generate an ICC profile [11] for this combination, a diffusor box was attached to the lamp in order to illuminate a test chart (GretagMacbeth Colorchecker DC) evenly. A set of images of the test chart with varying exposure time was captured. All images were converted to the camera’s native linear color space. Before creating ICC profiles with ProfileMaker 3.1.5 we applied a standard gamma correction ($\gamma = 2.2$) which improved the quality of the profiles. Out of the resulting set of profiles, the profile with largest gamut was selected for further use (which corresponded to correct exposure).

4.1. Data Acquisition

For each view we acquired a set of images $\{I_j\}$ with different exposure times T_j . HDR image generation follows standard procedures [3, 17]: We mark pixels in the input images as invalid for which at least one channel is overexposed or for which all channels are underexposed resp. below a certain threshold. All valid pixels are then converted to CIE XYZ color space [1] using the Little CMS color management engine [18] with high precision setting and relative colorimetric intent. This ensures that no gamut compression is applied and in-gamut colors are reproduced faithfully. The images are then converted to xyY color space in order to separate chromaticity and luminance channels. The luminance channel Y is used as a confidence measure and we define a weighting function $w(Y)$ that emphasizes reliable pixels, suppresses less reliable pixels, and is zero for invalid pixels.

We assume here that the reliability of the chromaticity information is correlated with the luminance channel (i.e., colors in bright areas are captured more reliably) and consequently define the weighting function only in terms of Y . Compared to previous work [6] where the original XYZ colors were used as parameters of the weighting function, the weighting function based on the luminance channel alone is not only adapted to the structure of the

xyY color space but also reflects the fact that chromaticity information in brighter areas seems to be better represented/converted by the profile mechanism.

The value of a pixel i in the final HDR image is then computed as

$$\begin{aligned} x_i &= \frac{\sum_j x_{i,j} w(Y_{i,j})}{\sum_j w(Y_{i,j})} \\ y_i &= \frac{\sum_j y_{i,j} w(Y_{i,j})}{\sum_j w(Y_{i,j})} \\ Y_i &= \frac{\sum_j Y_{i,j} T_j^{-1} w(Y_{i,j})}{\sum_j w(Y_{i,j})} \end{aligned} \quad (1)$$

The generated HDR image is converted back to CIE XYZ color space and used as input for BRDF generation.

5. Experimental Validation

Several objects – a clay model of a bird, a corroded bronze bust, and a porcelain caraffe – were acquired using the presented technique. We first compare the acquired appearance model with spectrophotometric measurements of the real objects in order to check color fidelity. In a second step, we render the model illuminated by a point light source and compare it to a photograph acquired under identical lighting conditions.

5.1. Spectrophotometric Measurements

We used a GretagMacbeth Eye-One spectrophotometer with a 45/0 measurement geometry to perform spectral measurements of the target objects at selected surface points. This proved to be difficult due to the curved surface geometry of the test objects and due to the bulky measurement head. In order to minimize errors, we selected near-flat regions and averaged multiple measurements. The BRDF data was generated by sampling at several locations and averaging the results. As the BRDFs are originally estimated up to a linear scale factor we determine an optimal scale factor per object based on the luminance channel and scale the BRDF data accordingly. Table 1 shows the resulting CIE XYZ values and x, y chromaticity coordinates. The difference between the measured and modeled values is given in ΔE (computed from the rescaled CIE XYZ values). All measurements were performed for illuminant D50 and 2 degree standard observer.

The results show that a ΔE of about 4 can be achieved for some of the materials. The orange and white surfaces of the bird show a much larger error. This is at least partly caused by the inhomogeneous color of the real model (often the yellow color of the base material shines through). Furthermore, the complex geometry of the orange and white parts of the bird and the white parts of the caraffe make sampling with the spectrophotometer difficult and less reliable.

5.2. Validation of the Complete Model

To evaluate the quality of the complete BRDF model (i.e., including the modeling of highlights), we acquired additional images of the objects illuminated by a point light source. The camera pose and the position of the light source were computed using the same calibration steps employed in the original BRDF acquisition process.

Figure 2 shows a side-by-side comparison of a photograph of the corroded bust and the acquired model (both images were converted to sRGB space for display using the color management work-flow). The false-color image shows the difference between them in ΔE .

The BRDF model is clearly able to capture not only the general appearance of the object but also many details: The overall color reproduction is quite good, highlights are preserved and have the correct shape. The face and neck region appear however to be desaturated (especially the dark green area on the bust's neck) while the helmet is reproduced correctly. The highlights are generally not crisp enough. These problems lead to an average ΔE of 9.8995 for which all non-black pixels with $\Delta E < 30$ were counted (we assume that larger differences are mainly caused by other problems – see below).

Some effects not related to the appearance model can be observed as well: By far the largest error occurs in regions where the shadow-boundaries are misaligned. This can be caused by a misregistration of light source and object as seen at the shadow of the nose. Another error source is inaccurate geometry which leads for example to the loss of detail in the hair region. The 3D geometry model used in this example was smoothed to remove noise and contains about 120k triangles which was not sufficient to represent small-scale details.

6. Calibrated Rendering and Output

Up to this point, we only validated renderings generated from the actual BRDF model that did not consider the properties of the output devices. We compare therefore in this section renderings displayed on a CRT screen and color printouts to the real objects under identical illumination conditions.

6.1. Rendering of Calibrated BRDF Models

For high quality output, our BRDF rendering algorithm computes high-dynamic range images with floating-point precision in real time. The rendered HDR images require clamping and scaling in order to generate a low-dynamic range version. The resulting images are then transformed by the ICC profile for low-dynamic display into a device dependent RGB or CMYK color space. Again, a suitable ICC profile needs to be generated for each output device. If no such profile is available, the sRGB color space [10] can be used as default.

	BRDF (CIE XYZ)	Spectrophotometer (CIE XYZ)	BRDF (xy chromaticity)	Spectrophotometer (xy chromaticity)	ΔE
Bird yellow	59.41, 61.23, 14.58	61.40, 61.51, 13.42	0.439 0.453	0.450, 0.451	4.158
Bird orange	42.28, 33.17, 8.02	43.73, 32.73, 5.95	0.507, 0.397	0.530, 0.397	7.923
Bird blue	34.99, 40.09, 46.65	32.74, 37.48, 48.16	0.287, 0.329	0.277, 0.317	4.283
Bird white	68.79, 73.35, 61.04	76.67, 78.75, 58.77	0.339, 0.361	0.358, 0.368	7.168
Caraffe blue	17.63, 18.67, 33.26	17.72, 19.34, 36.325	0.253, 0.268	0.242, 0.263	3.398
Caraffe white	57.16, 58.33, 52.11	54.00, 56.23, 43.27	0.341, 0.348	0.352, 0.366	7.003

Table 1: Comparison between the BRDF model in CIE XYZ color space and with spectrophotometric measurements. CIE XYZ values and x, y chromaticity coordinates are computed for illuminant D50 and 2 degree standard observer. The ΔE values are computed from the CIE XYZ values.




Figure 2: From left to right: Photograph of the corroded bust, rendered BRDF model, difference image (ΔE), diffuse texture computed from the BRDF model. The false-color difference image covers the range of $[0, 50]$; areas with larger difference (mostly misaligned shadow boundaries) were cropped.

An alternative approach to clamping and scaling would be to apply a tone-mapping operator [4] which compresses the dynamic range such that structures in dark and in bright regions still remain visible. To our knowledge, nobody so far has reported the application of a tone-mapping operator in the context of color management.

6.2. Validation Including Output Devices

Figure 4 shows two photographs of a comparison setup with a CRT screen on the left, the real object in the middle, and a printout on the right. (The exposure times of both halves differ in order to compensate for the different luminance levels of the CRT and the rest of the scene). The models were rendered under identical conditions (i.e., same camera view point and illumination) as the real object.

The rendered image is displayed on a calibrated CRT monitor. The printout was performed by a calibrated HP Color LaserJet 8550 color laser printer. The system utilizes the complete color management pipeline and no manual

color tuning was performed. Under these conditions we can expect the best possible color reproduction within the limits of the devices involved. Figure 4 shows that our BRDF measurement pipeline based on ICC color management succeeded to reproduce the appearance on different output devices with only small color deviations. Note that part of these are also caused by the digital camera used to capture the validation images.

The comparison between CRT display and printout shows also some of the limitations of the output devices: The dynamic range of the color laser printer is only 1:28 so that highlights appear very dim. In addition, Figure 3 shows a plot of the gamut (the range of reproducible colors) of the camera system (black), the monitor (green), and the color laser printer (red). The printer gamut is smallest so that colors are most likely wrong on the printout.

7. Discussion

The overhead of adding color management to the presented appearance acquisition system is limited as it es-




Figure 3: Gamut plot of the camera system (black), the monitor (green), and the color laser printer (red) used in Figure 4.

essentially replaces only the standard HDR computation routine. Compared to the benefit to reference the recovered appearance models to a standardized color space, it should be a worthwhile effort. This is especially evident in the comparison to spectrophotometric measurements in Section 5.1.

7.1. Accuracy of Color Management

The accuracy of the color management system both in color reproduction as well as in linearity is crucial for the approach. In the current system, we use therefore a rather aggressive weighting function $w(Y)$ to achieve reliable colors. Strongly overexposed or underexposed pixel values that are quite common in an HDR image series are problematic for the chosen color management setup and can bias the result. We expect however, that this problem could be addressed by optimizations in the profile generation phase.

7.2. CIE XYZ versus sRGB Color Space

The fact that CIE XYZ is one of the color spaces used in the PCS of ICC-based color management system makes it a convenient and obvious choice of working color space. Nevertheless, the data could also be transferred to (linear) sRGB color space at various stages of the processing pipeline.

We opted to perform all computations in CIE XYZ mainly due to conceptual simplicity but also due to its large gamut. Our goal is to create an accurate appearance model of an object for which we utilize the full gamut of our acquisition system that is larger than sRGB even if sRGB covers the gamut of most current display and output systems. Note, however, that the assumption of the BRDF model that there is no cross-talk between the three color




Figure 4: Photographic comparison of the appearance model with the original object. The images show a rendered image on a CRT screen (left), the original object under the same pose and illumination conditions (middle), and a color laser printer print-out of the rendered model (right). The luminance levels required a longer exposure times for the screen compared to the other parts. The lower image is composed from 3 individual images taken under identical conditions as the model was broken during the acquisition.

channels may be violated by the CIE XYZ color space with its overlapping spectra.

8. Conclusion

The presented image-based BRDF acquisition pipeline is long and complex: The target object is imaged with a digital camera under defined illumination converting from the physical (spectral) representation to tristimulus values. The color management system applies then a complex transform to convert to CIE XYZ color space. The parameters of the nonlinear Lafortune BRDF model are fitted and the model is evaluated resp. rendered for some illumination condition. Nevertheless, the result (which does not yet include conversion to an output or display color space) matches the reference data from spectrophotometer measurements quite closely. The presented pipeline allows for the first time to measure a standardized error between the measured appearance model and the original object. Our experiments show that color management techniques can be successfully applied to high-dynamic range imag-

ing, i.e., even beyond their original purpose. Especially image-based acquisition techniques can benefit from this development.

Acknowledgements

The authors would like to thank Villeroy & Boch for providing one of the measured objects. This work was funded in part by the DFG Schwerpunktprogramm V3D2, by the European Community within the scope of the ViHAP3D Project IST-2001-32641 “Virtual Heritage: High-Quality 3D Acquisition and Presentation”, and by the Max Planck Center for Visual Computing and Communication.

References

- [1] Commission Internationale de l'Eclairage. Colorimetry. Publication CIE No. 15.2, 1986.
- [2] Robert L. Cook and Kenneth E. Torrance. A reflectance model for computer graphics. In *Computer Graphics (Proceedings of SIGGRAPH 81)*, pages 307–316, August 1981.
- [3] P. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proc. SIGGRAPH 1997*, pages 369–378, August 1997.
- [4] K. Devlin, A. Chalmers, A. Wilkie, and W. Purgathofer. Tone Reproduction and Physically Based Spectral Rendering. In *Eurographics 2002: State of the Art Reports*, pages 101–123, 2002.
- [5] Andrew Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufmann, 1995.
- [6] M. Goesele, W. Heidrich, and H.-P. Seidel. Color Calibrated High Dynamic Range Imaging with ICC Profiles. In *Proc. 9th IS&T Color Imaging Conference*, 2001.
- [7] X. He, K. Torrance, F. Sillion, and D. Greenberg. A Comprehensive Physical Model for Light Reflection. In *Proc. SIGGRAPH 1991*, pages 175–186, July 1991.
- [8] R. W. G. Hunt. *Measuring Colour*. Fountain Press, 3rd edition, 1998.
- [9] Richard S. Hunter and Richard W. Harold. *The Measurement of Appearance*. Wiley, 1987.
- [10] Multimedia Systems and Equipment – Colour Measurement and Management – Part 2-1: Colour Management – Default RGB Colour Space – sRGB. Publication IEC 61966-2-1, 1999.
- [11] International Color Consortium. Specification ICC.1:1998-09, File Format for Color Profiles. available from <http://www.color.org>, 1998.
- [12] E. Lafortune, S. Foo, K. Torrance, and D. Greenberg. Non-Linear Approximation of Reflectance Functions. In *Proc. SIGGRAPH 1997*, pages 117–126, August 1997.
- [13] E. Lafortune, S.-C. Foo, K. Torrance, and D. Greenberg. Non-Linear Approximation of Reflectance Functions. In *Proceedings of SIGGRAPH 1997*, pages 117–126, August 1997.
- [14] H. P. A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based Reconstruction of Spatial Appearance and Geometric Detail. *ACM Trans. Graph.*, 22(2):234–257, 2003.
- [15] S. Marschner, S. Westin, E. Lafortune, K. Torrance, and D. Greenberg. Image-based BRDF Measurement Including Human Skin. In *Proc. 10th Eurographics Workshop on Rendering*, pages 131–144, June 1999.
- [16] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometrical Considerations and Nomenclature for Reflectance. National Bureau of Standards, 1977.
- [17] M. A. Robertson, S. Borman, and R. L. Stevenson. Dynamic Range Improvement Through Multiple Exposures. In *Proc. of the Int. Conf. on Image Processing (ICIP'99)*, pages 159–163. IEEE, October 1999.
- [18] M. M. Saguer. *Little CMS Engine – How to use the Engine in Your Application*, 2004. Available from <http://www.littlecms.com>.
- [19] G. Ward. Measuring and modeling anisotropic reflection. In *Proceedings of SIGGRAPH 1992*, pages 265–272, July 1992.

Biographies

Michael Goesele is a postdoctoral researcher in the computer graphics group of the Max-Planck-Institut für Informatik in Saarbrücken, Germany. He studied computer science at the University of Ulm and the University of North Carolina at Chapel Hill and received his doctorate from Saarland University in 2004. His research interests include acquisition techniques for computer graphics, image-based rendering, and color science.

Hendrik Lensch is a postdoctoral fellow in the computer graphics lab at Stanford University. He studied computer science at the University of Erlangen and the Royal Institute of Technology (KTH) in Stockholm. He worked as a research associate in the computer graphics group at the Max-Planck-Institut für Informatik in Saarbrücken, Germany, and received his doctorate from Saarland University in 2003. Lensch's research interests include acquisition techniques for computer graphics, appearance acquisition, and image-based rendering.

Hans-Peter Seidel is the scientific director and chair of the computer graphics group at the Max-Planck-Institut für Informatik and a professor of computer science at the University of Saarbrücken, Germany. He is an adjunct professor of computer science at the University of Erlangen, Germany, and at the University of Waterloo, Canada. Seidel's research interests include computer graphics, geometric modeling, high-quality rendering, and 3D image analysis and synthesis. He has published some 150 technical papers in the field and has lectured widely on these topics.