

Seeing Virtual Objects: Simulating Reflective Surfaces on Emissive Displays

Benjamin A. Darling and James A. Ferwerda; Munsell Color Science Laboratory, Rochester Institute of Technology, Rochester, NY

Abstract

In this paper, we work toward creating display systems that can present virtual proofs and replicas that behave like real physical surfaces in the actual lighting environment surrounding the display. The goal is to recreate the experience of directly viewing a reflective surface while using self-luminous display screens. We use a computational approach that models the physical light sources and actively tracks the observer to generate simulated reflections that are consistent with the observer's real viewing position and the real lighting present. The spectral composition of the ambient illumination is interactively sensed and the surface colors are calculated with a real-time multispectral rendering pipeline to allow the rendered model to automatically update its color with changes in the real lighting.

Introduction

Computers, presenting image content on display screens, are widely used in applications where the intention is to convey the appearance of object surfaces. Computer-based displays provide the means to rapidly evaluate the appearance of different options in proofing and material design applications, allow for the presentation of stimuli that would be difficult to construct for studies of material appearance, and support widespread dissemination of culturally significant objects, such as artwork, that otherwise may not be easily accessible. In these cases, a primary goal is for the appearance of the reproduction onscreen to represent the appearance of the physical object surface that is being portrayed. In this paper, we present a framework for displaying surfaces on self-luminous display screens that begins to recreate the experience of directly viewing a real physical reflective surface.

When viewing a real object surface, the patterns of light that ultimately reach the observer are a product of the light sources in the environment, the material properties of the surface, and the position of the observer. The appearance of the object changes with variations in the spectral composition or geometry of the incident light and also with changes in the geometric relationship between the illumination, surface, and observer. In contrast, display screens are typically viewed in a dim surround and present objects as part of a self-luminous image. The object in the image or interactive simulation is typically located in a virtual space separated from the viewer and has its own virtual illumination that is independent of the real illumination in the observer's environment.

In this paper, we work toward creating a display system (shown in Figure 1) that can present virtual proofs and replicas that behave like real physical surfaces in the actual lighting environment surrounding the display. This concept is achieved with traditional self-luminous display hardware using a computational approach that models the physical light sources and actively tracks the observer to generate simulated reflections that

are consistent with the observer's real viewing position and the real lighting present. This work is related to research on exact color soft proofing systems, where the light output of the screen is matched in absolute colorimetry to the diffuse surface colors of a physical hardcopy [1]. By incorporating aspects from realistic image synthesis [2], spatial augmented reality [3], and light-sensitive displays [4], our goal is to extend this paradigm to support a range of material surfaces. We incorporate texture and gloss properties by accounting for the geometric configuration of the illumination relative to the screen and using computer-graphics rendering methods to light the surface. For a virtual object to behave like a real surface, the virtual reflections need to update to remain consistent with changes in the geometry or the spectral composition of the illumination. The system accounts for the position of the viewer relative to the screen and the angle of the screen with respect to the real lights, which allows the luminance of diffuse reflections and patterns of specular reflections to change as they would for a real physical surface. Finally, by interactively monitoring the spectral composition of the ambient light and incorporating a real-time multispectral color rendering pipeline, the chromaticities of virtual surface reflections are able to change automatically to remain consistent with the real physical lighting present.

If a reproduction on an electronic display can sufficiently imitate the behavior of a real-world surface that is perceived as an illuminated reflective object, it can provide the advantages of a digital proof, while still retaining the object-mode appearance attributes of the physical surface it is attempting to portray. Though our current system still has limitations, we take initial steps toward this goal. In the following sections, we describe background and related work, present a methodology for modeling and recreating the appearance of reflective surfaces, and present an initial prototype system along with a description of its capabilities.



Figure 1. Left, a real physical painting and right, a virtual model created from captured data [5] that has been interactively rendered and displayed on a self-luminous display screen.

Background and Related Work

The creation of a system to reproduce the appearance of objects in real physical space has a basis in computer-based proofing, which provides methodologies for reproducing color and appearance attributes, as well as systems for spatially-augmented reality [3] that work to bridge the gap between the virtual and physical world.

Laihanen [6] developed an early system for exact soft proofing, attempting to produce an appearance match by reproducing the exact colorimetry (chromaticity and absolute luminance) of prints on a CRT screen. Recently, Hill [1] developed a display system for exact color proofing to allow for direct comparisons of hardcopy and soft-proof patches on an LCD screen in an illuminated light booth. Hill's system was calibrated to reproduce the exact colorimetry of a physical ColorChecker, using spectral data on the light sources and by adjusting the luminance output level of screen regions until they matched a physical mask placed on the screen. Research efforts have also investigated incorporating gloss properties into the proofing process. Gatt et al. [7] performed goniometric measurements to develop a predictive BRDF model for the gloss properties of printed materials. Patil et al. [8] developed a gloss soft-proofing system that generated simulated prints in a virtual environment by mapping the images to 3D planes and allowed the user to change the virtual viewpoint using QuickTime VR. The tangiBook system [9] used tracking information on the orientation of the screen and the real position of the observer to update the virtual reflections shown on a display screen, to provide natural forms of interactivity with virtual surfaces that had color, gloss and texture properties. The tangiBook system provided the capability to render virtual objects within different virtual lighting environments, but did not attempt to maintain spectral or geometric consistency with the real physical lighting environment.

The lighting-sensitive display [4] introduced the concept of illuminating virtual content on a display according to the lighting in the real environment surrounding the screen. This system interactively acquired an image map of the lighting in the environment with a camera and used this information to relight a static 3D scene using image-based rendering methods. Koike and Naemura [10] developed a display system capable of directionally modulating the light output using a lenticular array to simulate the view-dependent properties of surface BRDF. Fuchs et al. [11] explored different configurations of light field display systems capable of responding to spatial or directional incident light and outputting spatial or directional patterns of light.

In this work, we bring together concepts from colorimetric-based proofing systems, realistic image synthesis, and lighting sensitive displays in an attempt to create virtual objects under computer control that appear like the real physical objects they are portraying.

System Overview

The display system in this paper is designed to recreate colorimetrically the patterns of reflected light that a physical surface, positioned at the screen's location in the real lighting environment, would produce in the direction of the observer. Toward this goal, the system combines sensing and modeling of

the real-world lighting present, tracking technologies, and a multispectral real-time rendering engine to simulate the types of light-surface interactions that contribute to the appearance attributes of a surface (color, gloss, and texture). The results are displayed through a photometrically-calibrated screen so that the luminance and chromaticity of the emitted light can be matched to the light that would be reflected by an object at the screen's position.

Object Surface Modeling

The properties of the virtual surfaces, including diffuse color, are represented in manner that is independent of illumination so that the virtual output can be interactively updated as lighting changes, as a physical sample would change with the lighting.

Diffuse Color

Color calculations are performed using a multispectral factor methodology [12] by multiplying the coefficients of reflectance and illumination over a set of six optimized spectral channels developed for real-time multispectral rendering [13]. The diffuse color data for the virtual surfaces are represented in a multispectral form similar to reflectance factor, where the reflectance coefficient for each wavelength band varies between 0 and 1. To support spatially varying color for objects such as paintings or digital print proofs, six-channel reflectance data are maintained as two three-channel (RGB) floating point images.

Gloss

The gloss properties of the surface are represented using the specular reflectance parameters of the Ward BRDF model [14]. The Ward α (specular roughness) parameter is used to describe the width of the specular lobe. The parameter describing the magnitude of the specular reflectance, ρ_s , may be specified for each of the six multispectral channels to allow for spectrally-selective front surface reflection.

Small Scale Geometry and Texture

The system is intended to display surfaces with principally planar geometry, like the real screen has, so that the locations of virtual reflections are consistent with the physical location of light emitted from the screen. To support object surfaces with some texture or relief, small scale geometry (surface height $\ll$ viewing distance) is modeled as a dense height field relative to the physical plane of the screen. The effect of this geometry on the orientation of surface facets, relative to the planar screen surface, is represented with image-based normal maps [15]. The self-shadowing that results from higher surface points blocking light from reaching other surface points is represented with image-based horizon shadow maps [16].

Modeling the Light Booth Illumination

The system uses a model of the real-world illumination in the screen's environment to allow the virtual surface to be rendered in a manner consistent with a physical surface at the same position. For the lighting environment, we selected a light booth that has tungsten, fluorescent D50-simulating, and D65-simulating light sources.




Figure 2. Images captured of the spatial luminance variation for the three lighting options: (A) tungsten, (B) a fluorescent D50 simulator, and (C) a D65 simulator.

To account for the spectral composition of the illumination, the ambient light in the booth is continuously monitored while rendering using an Ocean Optics USB2000+ spectrometer, which samples the UV and visible spectrum at a refresh rate of approximately five spectra per second. These sub-nm spectral data are smoothed, resampled at 10 nm, and then converted to the six-channel multispectral representation with a $[36 \times 6]$ matrix representing the spectral sensitivity curves of the six optimized rendering primaries. The system has two spectral modes of operation. In the first mode, the direct spectral sensor mode, the continuously-updating six-channel illumination values from the sensor are used when rendering the color of virtual surfaces. The second mode is a classifier mode, where spectral data for each light source is pre-measured and the six-channel illumination values are stored. At rendering time, the spectral sensor is used to select between the three possible sources. Though the current implementation uses the spectral sensor in the classifier mode, this mode could potentially allow a simpler type of sensor to select between the light sources.

The geometric configuration and luminance of the light sources are also modeled. This information is used to estimate the incident irradiance levels at the virtual surface so that diffuse reflections can be rendered at the appropriate luminance levels and specular reflections can be rendered with correct spatial patterns.

The planar geometry for the ceiling of the light booth is specified in terms of a global coordinate system to describe its physical size and location. Though it has a diffuser, there is significant luminance variation across different regions of the booth light and this varies with the selection of the light source. The spatial luminance patterns for each of the light sources were captured in an offline process by taking a high-dynamic range image series of each light with a camera calibrated to estimate absolute luminance [17]. These images (shown in Figure 2) are geometrically warped to the planar geometry of the light booth ceiling to place the luminance patterns at their physical locations in the environment. The luminance pattern for the active light is chosen interactively from the three options using a classifier that selects the most likely light given the sensed spectral data.

To simplify rendering of diffuse reflections, a variance-minimizing median cut algorithm [18] is applied to the mapped luminance images to generate a set of representative point lights. These points store physical XYZ positions along with summed luminance values that represent different regions of the surface. The spatial luminance image is also maintained and is used when calculating specular reflections.

Screen and Observer Position

With the light sources modeled in physical space, it is possible to calculate geometrically consistent reflections if the position of the virtual surface on the screen and the positions of the observer are also known. As the observer moves relative to the display, the change in the real viewing direction is needed for updating the viewing angles in the BRDF calculation and determining the correct surface reflections at the observer's location. The user's head position is tracked using an IR-based system (NaturalPoint TrackIR) that provides the XYZ position of a set of markers worn by the user. To account for the position of the display screen, the XYZ position of a screen corner at a specified starting angle is measured. An accelerometer is used to interactively track the tilt angle of the display screen and its physical position is determined from the sensed orientation using a model of how the display rotates on its stand.

Rendering

The rendering component of the system uses the information on the screen and observer position along with the BRDF properties of the virtual surface to calculate the surface reflections for the modeled real-world lighting. These calculations are performed on the GPU and implemented using custom OpenGL shaders to allow them to be completed at interactive rates.

Diffuse reflections are calculated by iterating over the set of 32 pre-generated median cut lights for the currently active booth light source. The magnitude of the illuminance E at each screen location is calculated based on the physical light-to-screen-pixel distance and the orientation of the normal-mapped pixel relative to each light source:

$$E = \sum_{n=1}^{32} \frac{L_n (\mathbf{N}_{surf} \cdot \mathbf{I}_n) (\mathbf{N}_{light} \cdot -\mathbf{I}_n)}{d_n^2} \left(\frac{A}{P} \right) \quad (1)$$

where $\mathbf{I}_n$ is the surface-point to light-point unit vector for the n^{th} light point, d_n is the distance between these points in meters, L_n is the summed luminance stored in the n^{th} light point, $\mathbf{N}_{surf}$ is the

surface normal at the point on the virtual object, N_{light} is the normal to the plane of the area light, and the area term A/P is the physical area of the captured light surface divided by the number of pixels in the light image. This physical illuminance value is used to scale the product of the normalized six-channel multispectral illumination power distribution (S_j) and the multispectral diffuse reflectance ($\rho_{d,j}$) to determine the per-channel diffuse reflections on a luminance-based scale:

$$L_{out,j} = \frac{(\rho_{d,j})(S_j)}{\pi} E, \quad \text{for } j = 1 \text{ to } 6 \quad (2)$$

A similar calculation is performed to estimate the real diffuse reflection from the front surface of the display screen ($\rho_d = 0.002$). This estimated flare is subtracted from the virtual diffuse reflection to help mitigate the screen surface reflection present.

Specular reflections are rendered based on the specular term of the isotropic form of the Ward BRDF model [14] with the Dür [19] modification:

$$\rho_{brdf}(\theta_i, \phi_i, \theta_r, \phi_r) = \rho_s \frac{1}{\cos \theta_i \cos \theta_r} \frac{\exp(-\tan^2(\theta_h)/\alpha^2)}{4\pi\alpha^2} \quad (3)$$

where ρ_s is the specular reflectance parameter, α is a roughness parameter representing the width of the specular lobe, and θ_h is the angle between the surface normal and the half vector of the illumination and detection vectors. The specular reflections at each image point are estimated by real-time filtered importance sampling of the light source luminance image using an algorithm based on the method by Colbert et al. [20, 21] and the Ward model sampling equations described by Walter [22]. The six channel results from the diffuse and specular reflection calculations are summed and multiplied by a $[6 \times 3]$ matrix transform to CIE XYZ.

Display

The colorimetric results from the rendering engine provide the spatial pattern of light necessary to simulate the reflections from the virtual surface. The final stage is to translate these colorimetric values to the screen with a display characterization specified in absolute luminance levels. In order to match the luminance levels of reflections from real-world illuminated surfaces, we utilize a high-luminance display screen (EIZO RX220) capable of output levels as high as 900 cd/m² in a customizable mode, and 400 cd/m² in the luminance-stabilized mode. The display output was measured over the RGB and gray ramps using a Photo Research PR-655 spectroradiometer and the screen response was characterized using the model developed by Day et al. [23]. The determination of RGB digital counts for each pixel from the CIE XYZ output of the rendering engine is achieved at interactive rates using a custom GPU shader-based implementation of the inverse Day model.

In the current prototype, the screen characterization is based on measurements taken perpendicular to the screen. An IPS-based screen was selected in an effort to minimize the impact of viewing angle on the screen output. In preliminary measurements, it was found to maintain approximately 90% of the on-axis luminance at 15 degrees and 75% of the on-axis luminance at 25 degrees. A

characterization model that includes compensation for screen luminance changes at the tracked viewing position of the observer is in development.

System Capabilities and Results

Color Rendering and Display Evaluation

A measurement experiment was performed to evaluate the relative colorimetric accuracy of the six-channel rendering workflow and output through the display characterization. For the three types of illumination in the light booth, a simulated classic 24-patch ColorChecker was rendered. The light emitted by the display for each patch was measured with a Photo Research PR-655 spectroradiometer. All patches were measured at the same screen location by shifting the virtual model position on the screen. During the measurements, the system was operated in light classifier mode to maintain consistent virtual lighting data across all the patch measurements for a given light source. (In this mode, the spectral power distribution of each real light is pre-measured and stored as multispectral illumination values. The spectral sensor is used to select between the pre-measured light sources during rendering.)

For comparison to the patch data measured from the screen, we calculated the expected diffuse colors of the patches spectrally (10 nm intervals) for the measured spectral power distributions of the three light sources. Additionally, we generated simulated color results for the six-channel multispectral workflow. The multispectral results were calculated by first converting the measured spectral light distributions and measured surface reflectance curves to the six-channel representation. The light and surface data were multiplied on a per-channel basis and the result was converted to XYZ with a $[6 \times 3]$ transform matrix.

The baseline full-spectral calculation, the six channel multispectral calculation, and the values from real screen measurements were converted to CIELAB and compared using the CIEDE2000 color difference formula. Prior to the CIELAB conversion, the XYZ data from the different workflows were normalized to a relative scale by dividing the data by the Y value of the white ColorChecker patch in that workflow and multiplying by a factor of 1.13. For all three methods, the CIE XYZ value of the measured light source in that lighting condition was used as the starting white point. All data were then chromatically adapted to D65 using the chromatic adaptation transform from CIECAM02 so the magnitude of color error could be compared in a common space.

The mean, standard deviation, 90th percentile, and maximum CIEDE2000 color error between the baseline spectral calculation and the multispectral calculation for the ColorChecker are shown in the left portion of Table 1. The color error statistics between the

Table 1: Relative Color Error for ColorChecker Reproduction

		Six Channel Simulation				Screen Output			
		CIEDE2000				CIEDE2000			
		mean	std	90th	max	mean	std	90th	max
D65-Booth Light	(24)	0.4	0.3	0.8	1.7	0.8	0.4	1.4	1.6
D50-Booth Light	(24)	0.4	0.3	0.7	1.5	0.7	0.4	1.1	1.8
III-A-Booth Light	(24)	0.8	0.4	1.3	2.1	3.8	4.2	11.0	14.0
In-gamut patches	(18)	0.8	0.4	1.4	2.1	1.8	1.6	5.0	5.9

baseline spectral calculation and the measured color output from the display screen are shown in the right portion of the table.

The multispectral workflow provided reasonable color accuracy across lighting conditions, with mean CIEDE2000 below 1.0 for all three light sources, though there was a small increase in error for the illuminant A (tungsten) booth light. The measured screen output showed only slightly higher error than the simulated workflow for the D65 and D50 light booth conditions, but accuracy under the tungsten light was highly dependent on the patch. The maximum error increased to as high as 14.0, with a 90th percentile error of 11.0. The highest errors stemmed from screen gamut limitations. All the patches that exceeded a value of 6.0 were found to be out of gamut of the screen and so could not be physically produced. Five patches (orange, yellow, orange-yellow, yellow-green, red) fell outside the chromaticity triangle formed by the display primaries. Additionally, the black patch was out of gamut because the screen's minimum possible Z value (all channels at 0 digital counts) exceeded the small Z value needed to match a black patch under the tungsten illumination. For comparison, the results under illuminant A for the 18 in-gamut patches are also shown in Table 1. With the six out-of-gamut patches removed, the mean error decreased to 1.8, and maximum error decreased to 5.9.

System Capabilities

The system provides the capabilities to simulate diffuse color, surface texture, and gloss properties in a manner consistent with physical surfaces in the environment.




Figure 3. The surface of a virtual painting shown without its diffuse color to illustrate the shading and shadowing from the mesoscale texture. The rendered gray surface is surrounded by a real white mat, black mat, and black frame.




Figure 4. Sequence of screen captures illustrating increased shadowing in the rendered model when the screen is tilted downward and away from the real overhead light.



Figure 5. Left, the virtual reflection of the tungsten light booth source on a virtual curved surface. Right, the virtual reflection of the D65 source. Based on a spectral classifier, the system automatically identifies which light is present and switches the spatial light map used to calculate the reflections.

Texture

In Figure 3, a model of a virtual painting is shown with the diffuse color removed to illustrate the appearance of shading and shadowing from the surface texture. Using the orientation information provided by the accelerometer attached to the screen, the rendered shading and shadowing can automatically update when the surface is physically tilted. A sequence of images showing the increased shadowing as the surface is tilted downward and away from the booth light is shown in Figure 4. In order to illustrate the texture and shadow detail, each image in Figure 4 shows a small region of the surface that has been cropped from a direct screen capture. (Note that because they are screen captures, and not photographs from a fixed viewpoint, each captured image continues to show the view perpendicular to the screen as the display is tilted.)

Gloss

The system generates virtual reflections on surfaces that are consistent with the real spatial layout of the booth lights. The spatial luminance images depicted in Figure 2 are sampled according to the surface BRDF to create realistic specular reflections (shown in Figure 5). Using a spectral classifier to identify which lighting option is active, the system automatically switches between the spatial luminance pattern models so that the reflections change when the real light in the booth is changed.




Figure 6. A sequence of images illustrating the system automatically updating the color of a virtual ColorChecker (right side) to continue to match the color of the real ColorChecker (left side) as the light changes between the D65 light (left image), D50 light (center image) and tungsten light (right image). The images were purposely not color balanced (all were set to a CCT of 5500) in order to clearly show the changing illumination.

Color

The display system uses interactive sensing of the spectral composition of the light and the multispectral rendering pipeline to automatically update the color of virtual surfaces for changes in the illumination. This capability is illustrated in Figure 6, where a real ColorChecker is shown side-by-side with a virtual model on the display system. The displayed color onscreen updates as the lighting changes between the D65, D50, and illuminant A booth lights to maintain consistency with the physical ColorChecker. Additionally, with spectral sensing and a multispectral pipeline, the system has the capability to simulate metamerism, which would not be possible with a three-channel pipeline. A virtual target was created where each row contains a metameric pair of spectral reflectance curves. Each spectral pair was selected to create a metameric match under the tungsten booth light and to produce a color mismatch for the D65 simulator light. As shown in Figure 7, the samples appear to match under the tungsten light, but exhibit color differences under the D65 light.



Figure 7. A set of metameric samples displayed under the tungsten light (left) and the D65 simulator light (right). Each row of the target contains two virtual samples with spectra that were selected to provide a metameric match under the tungsten light and a color difference under the D65 booth light.

Conclusions and Future Work

In this paper, we have presented a framework for displaying virtual proofs and replicas on self-luminous display screens to simulate real-world reflective surfaces. This system supports color, gloss, and texture properties and provides methods for updating these attributes to maintain consistency with changes in the real-world illumination or the geometric relationship between the screen, observer, and lighting. Though the current system has limitations, we have taken initial steps towards the goal of presenting electronic objects on displays screens in a way that recreates the experience of directly viewing real surfaces.

While the current results are promising, there are limitations that suggest future work. Based on the results of the virtual ColorChecker measurements, it is evident that screen gamut limitations will present a challenge for exact color matching under illuminants that have a low correlated color temperature. A specialized wide-gamut screen may be necessary to match yellow or orange surfaces for these illumination conditions. There are also limitations imposed by the current need to capture the spatial luminance patterns and create a geometric model of the real lighting in an offline process. With the light booth used in the current prototype, the light sources have different spectral distributions so it is possible to automatically switch between the pre-generated spatial models with a spectral classifier approach based on the spectral sensor. The rendering engine is not limited to the light booth configuration and can handle a range of lighting options that can be represented as spatial luminance patterns mapped to geometric surfaces. However, if the lighting options do not have different spectral distributions, then the spectral classifier approach would not be able to automatically switch between them, and a different detection approach would be necessary to automatically update the spatial lighting model. An additional limitation of the current prototype is that only the direct illumination from the light sources in the viewing booth is considered in the lighting calculations. Indirect light reflected from booth surfaces is not included during rendering. Currently there are dark coverings on the side walls and part of the floor to minimize the amount of unmodeled light present. In future work, the luminance patterns from these surfaces of the booth could be captured and included as light sources in the lighting model. An additional limitation relates to the real physical reflection from the screen surface. In our light booth environment, where the light is

primarily from above, the reflection from the screen toward the observer is relatively small and diffuse, so it can be modeled by diffuse reflectance parameters and subtracted from the virtual image. Lighting geometries where significant light is directly in front of the screen would require other methods to account for unwanted physical screen reflections.

Computer-based reproductions presented on display screens provide a useful tool for simulating the appearance of virtual surfaces in a variety of applications including proofing, material design, and providing access to digital collections of artwork and other culturally-significant items. With continued development, a framework for simulating reflective surfaces has the potential to allow digital reproductions onscreen to appear like the real objects they are portraying, while providing the flexibility of computer-based methods.

Acknowledgements

This work was supported by NSF grant 0811680 made to James Ferwerda. We would like to thank Roy Berns for providing surface model data for the painting displayed in Figure 1.

References

- [1] B. Hill, Softproofing System for Accurate Colour Matching and Study of Observer Metamerism. Proc. CGIV '10, pg. 173. (2010).
- [2] D.P. Greenberg, K.E. Torrance, P. Shirley, J. Arvo, J.A. Ferwerda, S. Pattanaik, E. Lafortune, B. Walter, S. Foo, and B. Trumbore, A Framework for Realistic Image Synthesis, Proc. ACM SIGGRAPH '97, pg. 477. (1997).
- [3] R. Raskar, G. Welch, H. Fuchs, Spatially Augmented Reality, Proc. First International Workshop on Augmented Reality, pg. 1. (1998).
- [4] S. K. Nayar, P. N. Belhumeur, and T. E. Boult, "Lighting Sensitive Display," ACM Transactions on Graphics, 23(4), 963 (2004).
- [5] R.S. Berns and T. Chen, Practical Total Appearance Imaging of Paintings, Proc. Archiving 2012, pg. 162. (2012).
- [6] P. Laihanen, "Exact Soft Proofing," Journal of Imaging Science and Technology, 38(5), 432 (1994).
- [7] A. Gatt, S. Westland, and R. Bala, Increasing the Dimensionality of Soft Proofing: Gloss and Material, Proc. IS&T 14th Color Imaging Conference, pg. 292. (2006)
- [8] R. A. Patil, M.D. Fairchild, and G.M. Johnson, 3D Simulation of Prints for Improved Soft Proofing, Proc. IS&T 12th Color Imaging Conference, pg. 193. (2004).
- [9] B.A. Darling and J.A. Ferwerda, The tangiBook: A Tangible Display System for Direct Interaction with Virtual Surfaces, Proc. IS&T/SID 17th Color Imaging Conference, pg. 260. (2009).
- [10] T. Koike and T. Naemura, BRDF Display: Interactive View Dependent Texture Display using Integral Photography, Proc. ACM IPT/EDT '08, Article 6. (2008).
- [11] M. Fuchs, R. Raskar, H-P. Seidel, and H. Lensch, "Towards Passive 6D Reflectance Field Displays," ACM Transactions on Graphics (SIGGRAPH '08), 27(3). Article 58 (2008).
- [12] M.S. Drew and G.D. Finlayson, "Multispectral Processing Without Spectra," Jour. Opt. Soc. Am. A, 20(7), 1181 (2003).
- [13] B.A. Darling, J.A. Ferwerda, R.S. Berns, and T. Chen, Real-time Multispectral Rendering with Complex Illumination, Proc. IS&T/SID 19th Color and Imaging Conference, pg. 345. (2011).
- [14] G. J. Ward, "Measuring and Modeling Anisotropic Reflection," Computer Graphics, 26, 265 (1992).
- [15] M. Peercy, J. Airey, and B. Cabral, Efficient Bump Mapping Hardware, Proc. ACM SIGGRAPH '97, pg. 303 (1997).
- [16] P.-P. Sloan and M. F. Cohen, Interactive Horizon Mapping, Proc. Eurographics Workshop on Rendering, pg. 281. (2000).
- [17] P. Debevec and J. Malik, Recovering High Dynamic Range Radiance Maps from Photographs, Proc. ACM SIGGRAPH 1997, pg. 369. (1997).
- [18] K. Viriyothai and P. Debevec, Variance Minimization Light Probe Sampling, ACM SIGGRAPH Poster Proceedings, Article 92. (2009).
- [19] A. Dür, "An Improved Normalization for the Ward Reflectance Model," Journal of Graphics, GPU, and Game Tools, 11(1), 51 (2006).
- [20] M. Colbert, S. Pattanik, and J. Krivanek, "BRDF-Shop: An Artistic Tool for Creating Physically-Correct BRDFs," IEEE Computer Graphics and Applications, 26(1), 30 (2006).
- [21] M. Colbert and J. Krivanek, "GPU-Based Importance Sampling" In GPU Gems 3, H. Nguyen ed. (Addison-Wesley, Upper Saddle River, NJ, 2007) pg. 459.
- [22] B. Walter, Notes on the Ward BRDF, Tech. Rep. PCG-05-06, Cornell University, (2005).
- [23] E. Day, L. Taplin, and R. Berns, "Colorimetric Characterization of a Computer-Controlled Liquid Crystal Display," Color Research & Application, 29(5), 365 (2004).

Author Biography

Benjamin Darling is a Ph.D. candidate in the Color Science Graduate Program at the Rochester Institute of Technology. He received B.S. and M.S. degrees in Systems Engineering from the University of Virginia. Before entering the Ph.D. program at RIT in 2007, he was a research scientist at UVA working on direct brain interfaces and functional brain imaging.