Vision Research 109 (2015) 166–177

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

ImpastoR: A realistic surface display system

James A. Ferwerda

Chester F. Carlson Center for Imaging Science, Rochester Institute of Technology, United States

ARTICLE INFO

Article history:
Received 16 June 2014
Received in revised form 18 September 2014
Available online 15 November 2014

Keywords:
Material appearance
Display systems
Tangible interfaces

ABSTRACT

In this paper we describe our efforts to create a *surface display system* that produces realistic representations of real-world surfaces. Our system, based on a desktop PC with GPU hardware, LCD display, light and position sensors, and custom graphics software, supports the photometrically accurate and visually realistic real-time simulation and display of surfaces with complex color, gloss, and textural properties in real-world lighting environments, and allows users to interact with and manipulate these virtual surfaces as naturally as if they were real ones. We explain the system design, illustrate its capabilities, describe experiments we conducted to validate its fidelity, and discuss potential applications in material appearance research and computer-aided appearance design.

1. Introduction

Surfaces are everywhere, and it is through the interaction of light with surfaces that we perceive the properties of the world. We often create images of surfaces to document their properties, and it is self evident that images can serve quite well as surface representations. Nevertheless, the visual information we get from an image of a surface is not the same as the information provided by the surface itself.

We get a lot of visual information from interacting with surfaces, either through direct manipulation, or through observing a surface from different viewpoints, and conventional images do not support either of these behaviors. For this reason we have been working to develop novel *surface display systems* Darling (2013), that harness the power of digital modeling, computer graphics, and sensor technologies to produce new kinds of images that look and behave much more like the surfaces they represent.

Fig. 1 shows ImpastoR, the surface display system described in this paper. The system is displaying its rendering of an oil painting model next to the painting itself. Computer graphics hardware and custom software in the system allow the model to be rendered to the screen in real-time, with lighting consistent with the real lighting in the environment. Tracking systems change the rendering realistically as the display and observer move. The experience provided by the system is similar to that of holding and observing the real painting.

Our goal is to develop a surface display system that bridges the real and virtual worlds. To achieve this we have identified three requirements: (1) the images produced by the system must be *realistic*, accurately reproducing the appearance of real surfaces; (2) the images must be *responsive*, changing appearance appropriately with direct manipulation and changes in observer viewpoint; and (3) the images must be *situated*, appearing to be reflective surfaces in the observer's environment. In the following sections we first discuss related efforts in this area. Next, we describe the components of our system and illustrate its capabilities. We then describe experiments we conducted to validate the system's fidelity. Finally, we outline potential applications and discuss future work.

2. Related work

There has been strong interest in developing natural interfaces for interacting with virtual objects since the earliest days of computer graphics. Sutherland and colleagues, along with laying the foundations for 3D graphics, did pioneering work on developing head-mounted display systems and 3D input devices (Sutherland, 1968; Vickers, 1974) Significant advances have been made in both areas since that time (Burdea & Coiffet, 2003).

An alternative approach that supports direct interaction with virtual content is the CAVE system (Cruz-Neira, Sandin, & DeFanti, 1993). Users are surrounded by projection screens that display computer graphics renderings, and tracking devices provide support for direct interaction. Another projector-based approach is represented by the shaderLamps and iLamps systems developed by Raskar et al. (2001) and Raskar et al. (2003), as well as similar work by Bandyopadhyay, Raskar, and Fuchs (2001), Bimber and Raskar (2005) and (Konieczny & Meyer, 2006), where warped CG images are projected onto real three-dimensional objects to give the observer the experience of real objects with surface properties that can be changed under computer control.

E-mail address: jaf@cis.rit.edu

http://dx.doi.org/10.1016/j.visres.2014.10.016
0042-6989/

A third major approach for supporting direct interaction with virtual content incorporates spatially aware displays and tangible interfaces. In the Chameleon systems (Fitzmaurice, 1993; Tsang et al., 2002) computer monitors were fitted with six degree-of-freedom trackers to create spatially-situated displays whose virtual content changed depending on their real-world positions. Konieczny et al. (2005) combined trackers with a flexible handheld rear-projection screen and custom-built projector to allow users to navigate and view slices through 3D medical data by moving or bending the screen. In the Virtual Mirror system Francois and Kang (2003) incorporated a video camera pointed toward the user on a spatially-aware display, to create the impression of a mirror. This system also allowed interactive viewing of reflective daguerreotype images (Lazzari et al., 2002). A significant feature of all these systems is their support for direct interaction through tangible interfaces. The strength of tangible interfaces is that the affordances of the systems (lifting, tilting, and rotating the display) support rich and natural modes of interaction with the virtual content (Buxton, 2008; Ishii & Ullmer, 1997).

Our surface display system is also related to advanced displays that incorporate light-sensing and light modulating capabilities. In the "light-sensitive display" Nayar, Belhumer, and Boult (2004) introduced the concept of illuminating virtual content on a display based on the real lighting surrounding the screen. Yang et al. (2008) developed a system capable of directionally modulating the light output from the screen to simulate a 4D exitant light field. Koike and Naemura (2008) developed a similar system, based on the lenticular optics. Fuchs et al. (2008) developed a passive display system that extended beyond 4D to simulate six dimensions of spatially varying reflectance over a 2D surface. Wetzstein et al. (2011) have developed a light field display based on tomographically-optimized multi-layered LCDs and Lanman et al. (2011) has employed polarization fields to create an optically efficient dynamic light field display. A review of recent work in computational displays is given in Masia et al. (2013).

Since our goal is accurate simulation and display of surface appearance, our work also has commonalities with efforts in the area of soft-proofing for digital printing. Laihanen (1994) developed an early system for soft proofing, attempting to produce an appearance match by reproducing the exact chromaticity and luminance of prints on a CRT screen. Recently, Hill (2010) developed a display system for soft-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 target, 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, Westland, and Bala (2006) performed goniometric measurements to develop a predictive BRDF model for the gloss properties of printed materials. Patil, Fairchild, and Johnson (2004) 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.

Finally, we have been developing precursors to the current system for several years. In Darling and Ferwerda (2009) we described the tangiBook, a prototype surface display system that used a standard laptop computer with custom rendering software to produce dynamic images of surfaces that changed appearance with display and observer movement. In Darling and Ferwerda (2012) we described several of the core components of our system that provide accurate rendering and real-time interactivity. The further contributions of this paper are the in-depth description and illustration of the fully integrated system, discussion of physical experiments to validate the accuracy of the system, discussion of a psychophysical experiment to explore the extent to which the surface display produces the same visual response as a real surface, and presentation of applications of the surface display in both basic and applied appearance research.

Fig. 1. A real painting (left), and a representation rendered by our surface display system (right). The color, gloss, and texture properties of the painting are all rendered correctly for the real scene lighting, and the appearance of the surface changes appropriately with movement of the observer or display.

3. System design

Our surface display system is designed to reproduce the patterns of light reflected by real surfaces in a real lighting environment. Toward this goal we model the color, gloss and texture properties of surfaces; sense and model the spectral, spatial, and intensive properties of a real illumination environment; develop a real-time multispectral rendering engine that simulates light/surface interactions; display the results through a photometrically calibrated screen, and track the observer and display positions to provide natural modes of interaction. A diagram illustrating the components of the system is shown in Fig. 2.

Due to the complexity of the system and space limitations in this paper, the following sections summarize the components and methods used in the system. For interested readers, full conceptual background and implementation details are given in Darling (2013).

3.1. Surface modeling

3.1.1. Color

In the extreme, accurately modeling and rendering surface color requires representing the reflectance spectrum of the surface, the emission spectra of the light source(s) illuminating the surface, and the response properties of the display system. For the sake

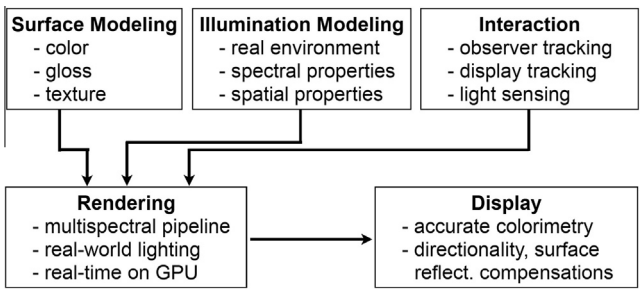

Fig. 2. Components of our surface display system. Elements of each component and their interactions are described in the text.

of performance, these full spectral calculations are typically reduced to products of normalized RGB values that are then scaled and sent as RGB digital counts to the display. While this approach is widespread, it can lead to significant errors in color rendering (Hall, 1989).

To preserve color accuracy while also supporting real-time graphics performance, we developed a multispectral rendering pipeline. Fig. 3 illustrates the pipeline, where surface and source spectra are each resampled into signals in six channels coded as ρ_d values in the Ward (1992) light reflection model with Dür (2006) corrections (hereafter referred to as the Ward–Dür model).

Since one of our interests was in accurately simulating artwork, our channel-sampling functions were derived from analysis of the spectral reflectance properties of an artist's material database and a set of standard illuminants. While these materials and illuminants had a wide variety of spectral properties, and we verified the accuracy of the sampling and rendering processes for surfaces of interest, six channels is not a magic number, and accurately rendering other materials and illuminants might require the use of more (or fewer) sampling functions with properties determined by the dataset.

The sampled spectral values were then used by a custom GPU-based shader we developed, to provide colorimetrically accurate renderings at real-time rates. Representative output of the system is shown in Fig. 4, where metameric color patches are correctly rendered as matching under tungsten light (source A) but not under daylight (source D65).

3.1.2. Gloss

Real surfaces often vary in gloss as well as color. We modeled surface gloss using the specular components of the Ward–Dür model, where ρ_s represents the specular reflectance, and α is a roughness parameter that controls the width of the specular lobe.

To model surfaces with spatially-varying reflectance properties, we stored the Ward–Dür parameters in image maps. Fig. 5 shows a painting with complex spatially-varying color and gloss properties, and three of the six image maps used to represent their Ward–Dür parameters. All parameters can be specified on a per-channel, per-pixel basis, offering great flexibility in surface reflectance modeling.

3.1.3. Texture

In addition to color and gloss, real surfaces often vary in height and orientation, which produce complex patterns of shading and shadowing that create rich visual textures.

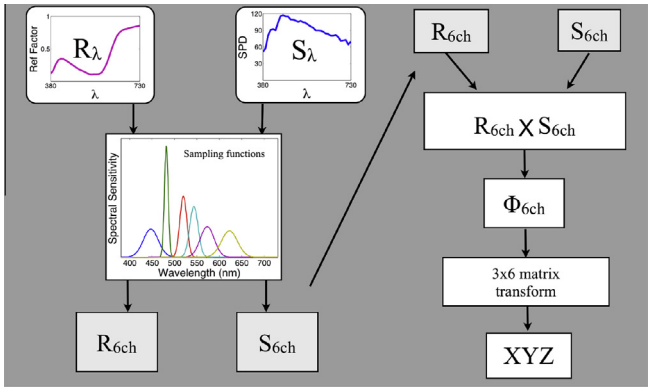

Fig. 3. Multispectral rendering pipeline. Surface reflectance spectra (R_λ) and light source emission (S_λ) spectra are multiplied by optimized sampling functions and integrated into six-channel representations (R_{6ch}, S_{6ch}). The products of these representations (Φ_{6ch}) are then transformed through a 3×6 matrix into CIE *XYZ* tristimulus values. This pipeline preserves color accuracy while providing real-time performance.

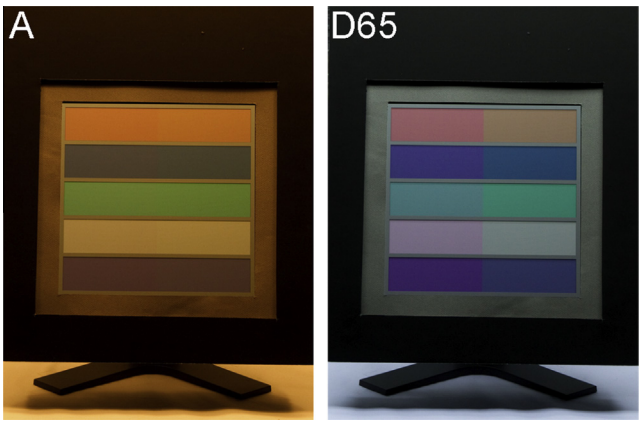

Fig. 4. Color accurate multispectral rendering. Color patches with different spectra that are metameric under source A are correctly rendered as matching, and are also correctly rendered as not matching under source D65. The Supplementary video shows that the lighting changes are sensed and color changes are rendered in real-time.

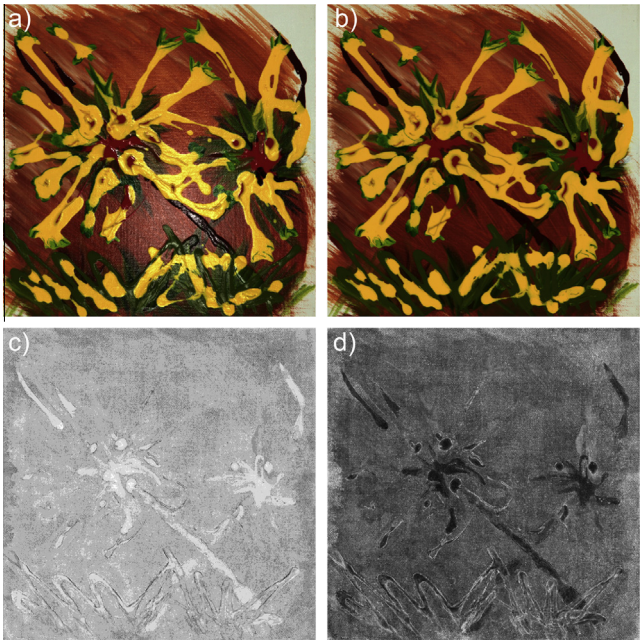

Fig. 5. System representation of surface properties (a) full rendering, (b) diffuse color (ρ_d), (c) specular reflectance map (ρ_s), and (d) specular roughness map (α). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We used surface normal maps to represent surface orientation and height maps to represent surface elevation at each surface point (Peercy, Airey, & Cabral, 1997). The normal maps were used during rendering to calculate surface shading, and the height maps were used to derive horizon maps to calculate shadowing effects (Max, 1988; Sloan & Cohen, 2000). Fig. 6 shows a painting with complex textural properties and a rendering produced by our system of the shading and shadowing effects caused by this texture.

3.2. Illumination modeling

Our system models a real-world lighting environment to allow rendered surfaces to appear consistent with real surfaces at the same position in the environment. For our lighting environment, we used a professional color evaluation booth (GTI ColorMatcher)

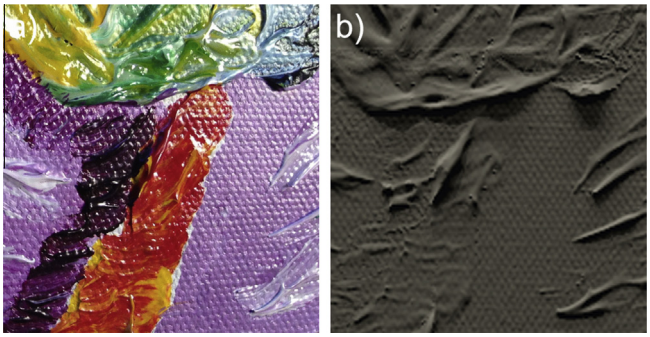

Fig. 6. System representation of texture. (a) Full rendering and (b) texture-only rendering showing shading and shadowing effects.

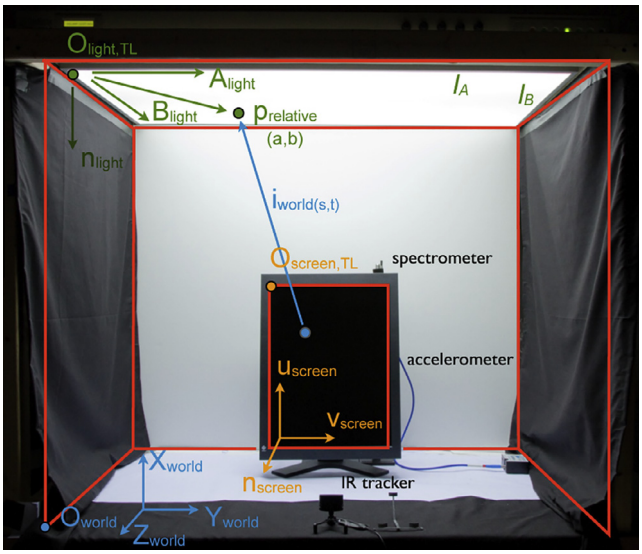

Fig. 7. Real-world lighting environment used in our system with overlays of world, screen and light source coordinate systems.

with switchable incandescent (CIE source A) and daylight fluorescent (CIE D50, D65) light sources (see Fig. 8). We modeled the spectral, spatial, and intensive properties of these sources.

3.2.1. Spectral properties

To account for the spectral properties of the sources, we monitored the light in the booth using an Ocean Optics USB2000+ spectrometer, which sampled the visible spectrum at 5 Hz. The sensor was attached to a pole behind the display screen and directed toward the top of the booth. The 10 nm spectral data from the sensor was resampled into the multispectral representation described earlier and used in the rendering process. The real-time spectral sensing allowed the system to respond automatically when sources were switched.

3.2.2. Geometry and luminance

We also modeled the spatial luminance distributions of the light sources. We captured these properties in an offline process by taking high-dynamic range images of each source with a calibrated digital camera. We then mapped these images to the ceiling of a geometric model of the booth and used them as illumination maps when rendering. During the rendering process the illumination map for the active source was chosen automatically using a classifier that selected the appropriate map given the sensed spectral data.

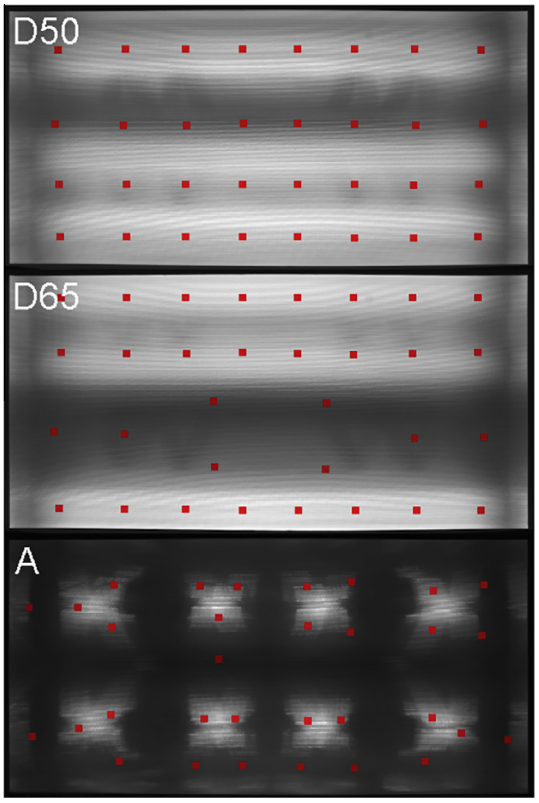

Fig. 8. Sources available in the light booth with overlays showing the point source approximations produced by the median cut algorithm (see Section 3.5).

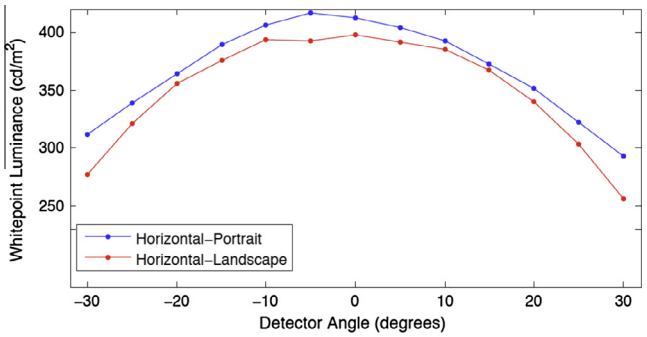

Fig. 9. Luminance fall-off for our LCD display as a function of horizontal detector angle. The two curves plot the results for portrait and landscape modes. These effects were compensated for during rendering.

3.3. Interaction tracking

3.3.1. Observer tracking

We used a NaturalPoint TrackIR infra-red (IR) tracking system to track the position of an observer using the system. The observer wore a headband with a pattern of IR reflective dots attached. The dots were illuminated by an IR source, and the pattern was captured by an IR camera. The tracker's sampling rate was 120 Hz, and its angular resolution was approximately 0.0065° within a 46° field of view. A software API used this information to track the position of the observer relative to the camera. We used this tracking data with offsets to calculate the observer's position relative to the lighting environment. The geometric relations used to determine the observer's position are illustrated in Figs. 7 and 10.

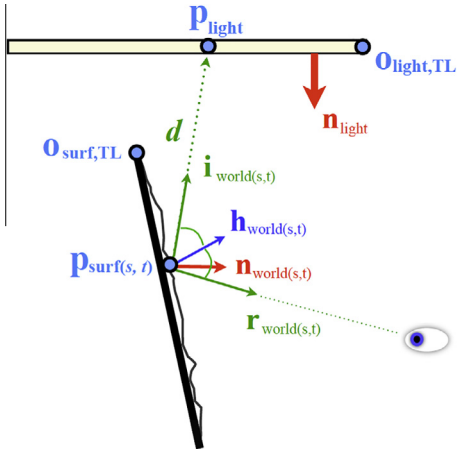

Fig. 10. Direction vectors used to calculate surface reflections for the real-world light source.

3.3.2. Display tracking

We also tracked the display screen to determine its orientation and the resulting positions of screen points relative to the lighting environment. To accomplish this we used an Action XL FP100W tri-axial accelerometer that was affixed to the back of the display. The accelerometer's sampling rate was 50 Hz and its sensitivity to changes in orientation was approximately 1.4° in each axis. We developed a model to account for the positions of screen points based on the accelerometer rotation and offsets from the center of rotation of the display. The relations used to determine the screen positions are illustrated in Figs. 7 and 10.

System latency was not measured explicitly, but response lag was very small, providing the experience of real-time interaction.

3.4. Display characterization

The display is one of the most critical components in the system. Several properties of the display impact the system's ability to accurately reproduce surface appearance. First the display's luminance range and color gamut must be large enough to reproduce the range of light produced by a real surface in the lighting environment. This output volume needs to be sufficient to include both surface highlights and shadows as well as the chromatic properties surface points. Next, the directionality of the display's output is also an important factor in determining how well the display can reproduce surface reflections, since the display will be manipulated with respect to the illumination and viewed from different directions. Finally, the front surface reflectance of the display is important because reflections from the screen will add to any light produced by the screen, and will need to be compensated for to accurately reproduce the appearance of a real surface. Our efforts to characterize and calibrate the display used in our system are described in the following paragraphs.

3.4.1. Colorimetry

To reproduce the luminance levels and dynamic ranges of reflections from real-world surfaces we used an EIZO RX220 high intensity medical LCD display in its 400 cd/m^2 luminance stabilized mode, which we further restricted to a peak of 320 cd/m^2 to allow for directionality compensation. We measured the black level of the display at 0.73 cd/m^2 yielding a dynamic range of 438:1.

We characterized the display colorimetrically using the method described by Day, Taplin, and Berns (2004), by measuring the CIE *XYZ* tristimulus values of the display's white point, black point, and RGB primary and neutral ramps using a PhotoResearch PR-655 spectrophotometer. We then used these data to linearize the display's response by creating 11-bit look-up tables (LUTS) to convert from RGB digital counts to normalized RGB scalars and derive a matrix to transform from displayed RGB scalars to *XYZ* tristimulus values. The form of this forward transformation is shown in Eq. (1). To go from calculated *XYZ* values to RGB digital counts for output we inverted the transformation and the LUTS.

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \left(L_{factor(\theta,\phi)} \begin{bmatrix} X_{R\max} - X_K & X_{G\max} - X_K & X_{B\max} - X_K \\ Y_{R\max} - Y_K & Y_{G\max} - Y_K & Y_{B\max} - Y_K \\ Z_{R\max} - Z_K & Z_{G\max} - Z_K & Z_{B\max} - Z_K \end{bmatrix} \right) \begin{bmatrix} R \\ G \\ B \end{bmatrix} + \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{K(\theta,\phi)} \tag{1}$$

3.4.2. Directionality

The light emitted from an LCD display typically varies as function of viewing angle. We selected a display with good angular uniformity to minimize the correction required, but given the freedom of viewing and manipulation possible in our system, we still needed to compensate for directional effects. To do this we measured the *XYZ* tristimulus values of the display's white point, black point, gray levels, and maximum RGB primary output at 5° intervals over a ±30° range from the display normal in both the landscape and portrait orientations. We found no significant changes in output chromaticity over the range, but as shown in Fig. 9, luminance declined at wider viewing angles. To compensate for this, we fit the data with an elliptical paraboloid parameterized by zenith θ (deviation from screen normal) and azimuth ϕ (direction along screen surface) angles, and used this model to derive a scaling factor $L_{factor(\theta,\phi)}$ that we included in the display characterization shown in Eq. (1). Using this formulation, the luminances of colors viewed off-axis are increased to compensate for the falloff. As discussed earlier this required restricting the on-axis peak luminance output to 320 cd/m^2 down from 400 cd/m^2 to provide headroom for the compensation.

3.4.3. Front surface reflectance

Since the display is used in an illuminated environment, there will be some light reflected from the display screen that may cause unwanted reflections. We chose a display with low front-surface reflectance, but we also measured screen reflectance using a calibrated light source and the PR-655 mounted on a manual gonio-arm. We found that peak screen reflectance at the specular angle was 1.1% and that reflectance fell off rapidly at off-specular angles and was negligible beyond 10° off-specular. We used this information during the rendering process to compensate for the display's front surface reflections as a function of display/light source orientation and observer viewing angle.

3.5. Rendering

The rendering component of the system uses real-time information on the screen and observer positions along with the reflectance properties of the modeled surface to calculate surface reflections for the modeled real-world lighting. Rendering was performed on a desktop gaming PC running Windows 7 with an NVidia GTX 670 graphics card, and is implemented using custom OpenGL shaders that provide output at interactive rates. The following sections survey the elements of the rendering process.

3.5.1. Diffuse color

Diffuse surface reflections are calculated by iterating over the set of 32 pre-generated median cut lights for the currently active booth light source (see Fig. 8). The magnitude of the illuminance

E at each screen location is calculated from the physical light-to-screen-pixel distance and the orientation of the normal-mapped pixel relative to each light source

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

where $\mathbf{I}_n$ is the surface-point to light-point unit vector for the nth light point, d_n is the distance between these points in meters, L_n is the summed luminance stored in the nth light point, $\mathbf{N}_{surf}$ is the surface normal at the point on the virtual object, $\mathbf{N}_{light}$ 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 multispectral illumination representation S_j and the multispectral diffuse reflectance $\rho_{d,j}$ to determine the per-channel diffuse reflections on a luminance-based scale:

$$L_{out,j} = \frac{(\rho_{d,j})(S_j)}{\pi}E, \quad \text{for } j = 1 \text{ to } 6 \quad (3)$$

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

3.5.2. Specular reflections

Specular reflections are rendered using the specular term of the Ward–Dür model:

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

where ρ_s is the specular reflectance, 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.

One of our principal challenges in rendering specular reflections was developing a workflow that could produce accurate renderings at real-time rates. While median cut point lighting (Viriyothai & Debevec, 2009) provides good accuracy and high performance when rendering low gloss surfaces, renderings of high gloss surfaces can show artifacts. Conversely, filtered importance sampling (Colbert, Premože, & François, 2010) often shows artifacts when rendering low gloss surfaces, but excels in accuracy and performance in the high gloss range. To produce accurate renderings, at real-time rates, for surfaces with a wide range of gloss properties, we developed a hybrid approach that automatically switched between the two techniques on a per-pixel basis, using a custom error minimization metric.

3.5.3. Shadowing

Since the system was designed to simulate the appearance of surfaces with texture or "impasto", surface relief was modeled using image-based height maps. Shadowing effects were represented using horizon shadow maps derived from the height maps (Max, 1988; Sloan & Cohen, 2000). For each surface point, a horizon map stored the critical elevation angle at which light would be blocked from reaching that point from a given azimuthal direction. We divided azimuth into 18 segments of 20° each. Employing three channels per image, the 18 horizon maps were stored as six floating-point images. In the rendering shader, the horizon map corresponding to the calculated azimuth angle to a light source was selected on a per-pixel basis. If the elevation angle to a source was below the critical angle stored in the map, that source was not included in the rendering calculation for that pixel, producing the shadowing effect.

3.5.4. Lighting coordinate system

In the system, the geometric relationships between the real-world light source, surface/display, and observer are considered, and the rendering equation is evaluated with respect to the real-world coordinate system shown in Fig. 7. When rendering with filtered importance sampling, the illumination direction is determined using the Ward–Dür model half-vector (calculated from the accelerometer orientation data, modified by the surface normal and the real viewing vector of the observer derived from the IR tracker data). The four direction vectors required to evaluate surface reflections ($\mathbf{r}_{world}$, $\mathbf{h}_{world}$, $\mathbf{n}_{world}$ and $\mathbf{i}_{world}$) are shown in Fig. 10. Using these vectors, it is possible to calculate the physical position on the light source that will be intersected by the illumination vector, so that the corresponding lookup into the light source illumination map can be performed.

The position on the light source plane is determined by ray-polygon intersection. The distance along the illumination vector from the object surface position (s,t) to the light plane is calculated by:

$$d = \frac{(\mathbf{o}_{TL,light} - \mathbf{p}_{surf(s,t)}) \cdot (\mathbf{n}_{light})}{(\mathbf{i}_{world(s,t)}) \cdot (\mathbf{n}_{light})} \quad (5)$$

where $\mathbf{p}_{surf(s,t)}$ is the physical position of the object surface point, and the light plane is specified by the plane normal $\mathbf{n}_{light}$ and a point on the plane $\mathbf{o}_{TL,light}$. The light intersection point $\mathbf{p}_{light}$ is found by moving a distance d along the illumination unit vector:

$$\mathbf{p} = d(\mathbf{i}_{world(s,t)}) + \mathbf{p}_{surf(s,t)} \quad (6)$$

Fig. 7 illustrates that to determine the corresponding pixel in the light source luminance image, the 2D parameterized position on the light source is found by determining the 3D position of this point relative to the origin of the light coordinate system:

$$\mathbf{p}_{relative} = (\mathbf{p}_{light} - \mathbf{o}_{TL,light}) \quad (7)$$

and then a pseudo-inverse calculation is used to determine the unit pixel coordinates (a, b) for the intersection:

$$\begin{bmatrix} a \\ b \end{bmatrix} = pinv\left(\begin{bmatrix} l_A \cdot A_x & l_B \cdot B_x \\ l_A \cdot A_y & l_B \cdot B_y \\ l_A \cdot A_z & l_B \cdot B_z \end{bmatrix}\right)\mathbf{p}_{relative} \quad (8)$$

where $\mathbf{A} = (A_x, A_y, A_z)$ and $\mathbf{B} = (B_x, B_y, B_z)$ are the unit left-to-right and top-to-bottom directions of the light source rectangle in world space, and l_A and l_B are the physical dimensions of the light source, as illustrated in Fig. 7. These (a, b) values serve as the texture lookup coordinates to index the light source luminance image to calculate surface reflections.

4. System capabilities

The system described above meets the three requirements we set out for surface display systems. First, the images produced by the system are *realistic*, and faithfully represent color, gloss and texture of modeled surfaces. Second, the images are *responsive*, and change appearance appropriately with direct manipulation and changes in user viewpoint. And third, the images are *situated* with respect to the scene illumination and the observer, and appear to be an integral part of the user's environment. In the following sections we illustrate the system's capabilities.

4.1. Color

Fig. 11 illustrates the system's color simulation capabilities. Each sub-image shows a real Macbeth ColorChecker chart on the left and the chart image produced by the system on the right. Note first that the color rendering is realistic due to the multispectral

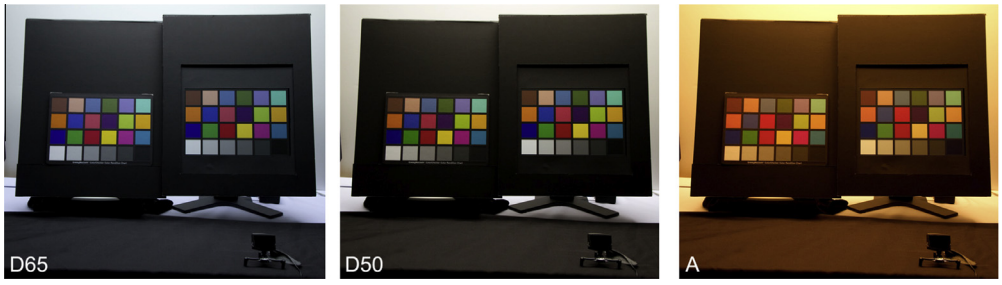

Fig. 11. Simulating color. In each panel the real ColorChecker chart is on the left, the image produced by the system is on the right. Note that the simulated colors change appropriately with respect to the light sources. The Supplementary video shows that the changes occur automatically and in real-time as the sources are switched.

color workflow used in the system. Note also that because the system can sense the spectrum of the light source, the colors in the simulated chart change appropriately for the different sources. The Supplementary video shows that the changes occur automatically and in real-time as the sources are switched.

4.2. Gloss

Fig. 12 shows the system's ability to simulate the appearance of glossy surfaces. Note that the hues of the neutral surfaces change appropriately for the different light sources, and that the reflections are correct for the different patterns of lights in the sources. Also, as shown in the Supplementary video, because of the system's sensors, the locations of the surface highlights also change appropriately and in real-time with observer and display movement.

4.3. Texture

Fig. 13 illustrates the system's ability to simulate the shading and shadowing effects produced by surface textures. The images show renderings of the canvas and brushstroke textures of a laser-scanned oil painting. Note that the surfaces show appropriate surface shading effects, with surfaces elements oriented toward the light source appearing brighter than those facing in other directions. Note also that regions that are occluded from direct illumination are appropriately shadowed. The two images show the effects of rotating the display monitor. Note that the shading and shadowing effects are different at the two orientations, correctly simulating the interactions of illumination and surface geometry. As shown in the Supplementary video, because of the system's sensors, these effects vary in real-time with changes in user viewpoint and manipulation of the display.

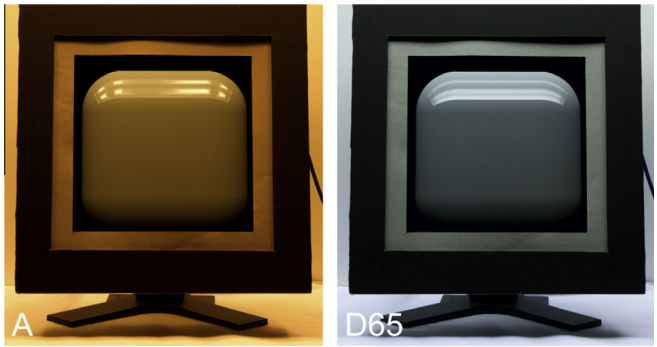

Fig. 12. Simulating gloss. Note that the colors and patterns of the highlights are correct with respect to the different light sources. The Supplementary video shows that the positions of the highlights change appropriately and in real-time with user or display movement. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

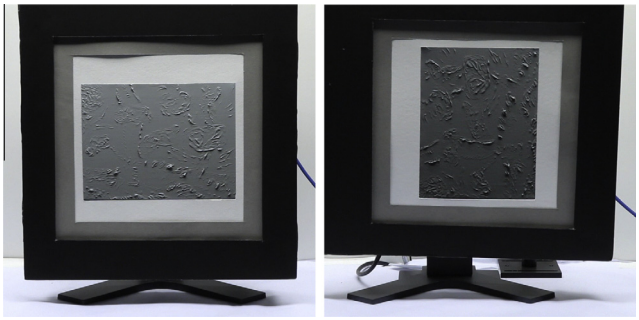

Fig. 13. Simulating texture. Note the surface shading and shadowing effects shown in the two panels are correct with respect to the surface's orientation to the light source. As shown in the Supplementary video, these changes happen automatically and in real-time as the user rotates the display.

4.4. Total appearance

Finally, Fig. 1 shows the system's capabilities brought together in one simulation. On the left is a real oil painting with complex, spatially-varying color, gloss, and textural properties. On the right is the rendering produced by the system. As shown in the Supplementary video, the rendering is dynamic and interactive, and responds as the real painting would to direct manipulation, changes in viewpoint, and changes in illumination.

5. System validation

While the results shown in the previous section are visually compelling, we want to produce simulated surfaces that are measurably accurate. In the following sections we describe two measurement experiments we conducted to assess the color and gloss reproduction accuracy of our system.

In the first experiment we evaluated the color reproduction accuracy of the system by simulating and displaying a Macbeth ColorChecker chart and comparing spectroradiometric measurements of the displayed colors with measurements of a real chart. In the second experiment we evaluated the gloss reproduction accuracy of the system by simulating and displaying a set of gloss samples and comparing goniometric measurements of the displayed samples with those of the real samples.

5.1. Color reproduction

We created a geometric model of a ColorChecker chart with the same dimensions as a real chart. We then measured the spectral reflectance of each patch in a real chart using an XRite Eye-One Pro contact spectrophotometer (to capture spectral reflectance independent of source irradiance). These spectral reflectance data

were then resampled into our multispectral representation and used to define the ρ_d reflectance values of patches in the model. The spectral power distributions of the three light sources in our booth were measured using a PhotoResearch PR-655 tele-spectroradiometer and a PTFE white reflector (to capture spectral irradiance independent of surface reflectance) Similar to above, these data were then resampled into our multispectral representation. The chart model was then rendered using each light source and the resulting images were displayed. For each of the chart/source renderings, the light spectra emitted by the displayed patches were measured with the PR-655.

XYZ tristimulus values for each of the patches were calculated using the CIE 2° standard observer color matching functions. In addition a set of *XYZ* values for the pre-display simulations were calculated to allow analysis of the display's contributions to any color differences. The sets of *XYZ* values were converted to CIELAB coordinates and compared using the CIEDE2000 color difference formula.

The mean, 90th percentile and maximum CIEDE2000 color differences between the real and simulated charts are shown in the middle column of Table 1. The color differences between the real and displayed charts are shown in the right portion of the table. As shown on the left, the six-channel multispectral workflow provided a very good level of color accuracy across lighting conditions. The mean CIEDE2000 for each of the D50 and D65 sources was 0.4. While there was increased error for tungsten source A, the mean CIEDE2000 value was still only 0.8. For reference, typical threshold values for perceived color differences have been found to be between 2.0 and 3.0 (Mahy, Van Eycken, & Oosterlinck, 1994).

As summarized in the right column of Table 1, the displayed output showed only slightly higher color differences than the simulations for the D65 and D50 sources, but the differences in increased dramatically in the displayed image simulated using source A. Here the maximum error was 14.0, with a 90th percentile error of 11.0. These errors can be attributed to screen gamut limitations. Under source A, six patches (orange, yellow, orange-yellow, yellow-green, red, and black) fell outside of the display's gamut and could not be physically reproduced by the display As the table shows, with these six out-of-gamut patches removed, the mean error decreased to 1.8, and maximum error decreased to 5.9. It should be noted that due to the effects of the incandescent illuminant, the chromaticities of these patches are quite extreme and would likely be out-of-gamut for many display systems.

Overall the results indicate that the color reproduction capabilities of our system are very good, but that the gamut limitations of our display may introduce visible errors for extreme out-of-gamut colors. This is a problem for any color-critical application of rendering and display technology, and can be minimized by tuning the color representation to the dataset of interest and confirming that the display gamut is sufficient to reproduce the colors in that set.

Table 1
Color validation. Color differences between a real ColorChecker chart and our system's reproductions.

	Six channel simulation CIEDE2000				Screen output CIEDE2000			
	Mean	Std.	90th	Max.	Mean	Std.	90th	Max.
D65 (24)	0.4	0.3	0.8	1.7	0.8	0.4	1.4	1.6
D50 (24)	0.4	0.3	0.7	1.5	0.7	0.4	1.1	1.8
IIIA (24)	0.8	0.4	1.3	2.1	3.8	4.2	11.0	14.0
IIIA (18)	0.8	0.4	1.4	2.1	1.8	1.6	5.0	5.9

5.2. Gloss reproduction

We also evaluated our system's ability to reproduce surface gloss by measuring and simulating set of gloss samples. The samples were 4 cm × 4 cm glass plates that were painted with glossy gray latex paints made by mixing high gloss and flat paints in different ratios. Nine samples were created with high gloss percentages of (100%, 98%, and incremental 5% mixes between 95% and 60%). The sample surfaces were generally flat, uniform, and without significant texture.

The diffuse reflectance values of the physical samples were measured with an X-Rite Eye-One Pro spectrophotometer, which uses a 45/0 illumination/detector measurement geometry. These spectral measurements were resampled into ρ_d values in our multispectral representation.

The gloss properties of the samples were estimated using two methods. For the higher gloss samples we imaged a line source and measured the resulting light distributions with a calibrated digital camera. For the lower gloss samples we measured the BRDFs with a Murakami GCMS-10X goniospectrophotometer. The image-based method provided high-resolution angular measurement for samples with narrow specular lobes, and the goniometer-based method allowed broader lobes to be measured over a wide angular range. Ward–Dür ρ_s and α parameters were fit to the measured distributions, and the models of the sample plates were simulated and displayed by the system.

To validate the fidelity of the simulations, we constructed a custom gonio-spectroradiometer to allow direct measurement of the real and simulated samples. The device, shown in Fig. 14, consisted of a calibrated light source (Mille Luce M1000 lamp, fiber optic cable, integrating sphere, filters to achieve a 4700 K correlated color temperature, and a 2.3 cm × 3.5 cm aperture) and a manual gonio-arm on which we mounted the PR-655. Either a sample holder or the display could be mounted at the center-of-rotation of the device. Using this device we measured the luminance produced by the real and simulated samples between 6° and 30° from the surface normal. Samples were taken at 0.5° intervals near the 7.5° specular angle and at progressively larger intervals at larger off-specular angles. The light source was on for the real samples and off for the simulated samples.

The absolute luminance values produced as a function of angle for the real and simulated samples are shown in Fig. 15. Although

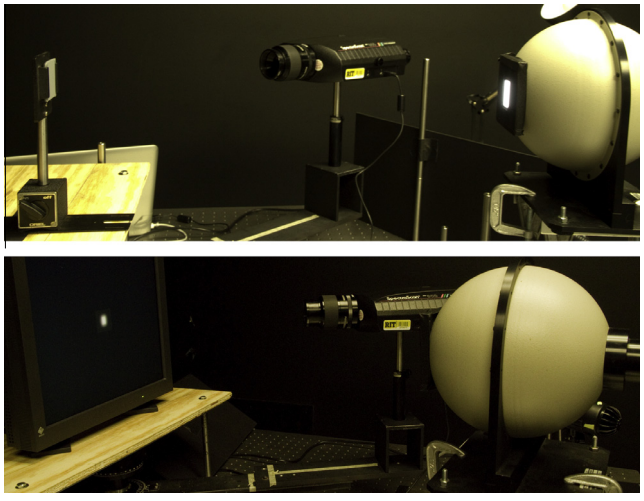

Fig. 14. Measurement setup for our gloss validation experiments. The top panel shows the arrangement used to measure the real gloss samples. The bottom panel shows the arrangement used to measure the displayed reproductions.

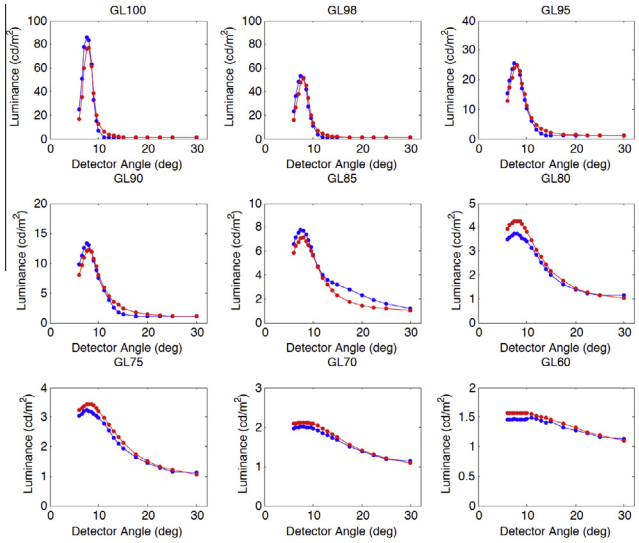

Fig. 15. Gloss validation Goniometric measurements of the luminance from real physical gloss samples (red lines) and our system's reproductions (blue lines). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

there are some deviations, in general the reflectance functions are similar in magnitude and shape. In the higher gloss samples (GL 85–100), modeled using the line source, the peak values of the simulated samples were marginally higher than the real samples. Conversely in the lower gloss samples (gloss level 60–80) modeled using the goniospectrophotometer, the peak values of the simulated samples were marginally lower than the real samples. The two mid-range gloss samples (GL 80 and 85) show the biggest differences and these were also the most difficult samples to measure and model, since the widths of their specular lobes fell between the measurement capabilities of the line source and goniospectrophotometer methods.

In addition to the physical measurements, high dynamic range images of the real and simulated samples were created using a calibrated digital camera mounted to the gonio-arm and set at the specular angle. The images of real and simulated samples GL100. GL95 and GL85 are shown in Fig. 16.

Overall the results show that our system can reproduce surface gloss with good fidelity as long as the surfaces reflectance properties can be accurately measured, and then described by the Ward–Dür model. While this covers a wide and useful range of materials, reproducing more complex materials may require more complex measurement methods and models. In addition, while the current system can accurately simulate gloss appearance in the lightbooth environment, generalizing the system to more varied or extreme lighting environments might require the use of a high dynamic range display, and evaluation of whether that display's "gloss gamut" can reproduce the intensities and contrasts produced by a given combination of surface reflectance and lighting.

6. Applications

While the primary focus of this paper has been on the challenges of developing a surface display system, the primary *motivation* for this work is to develop a tool to facilitate the study of surface and material perception. Although our research in this area is just beginning, in the following sections we outline how we are starting use the surface display to address questions that would be difficult to investigate by other means.

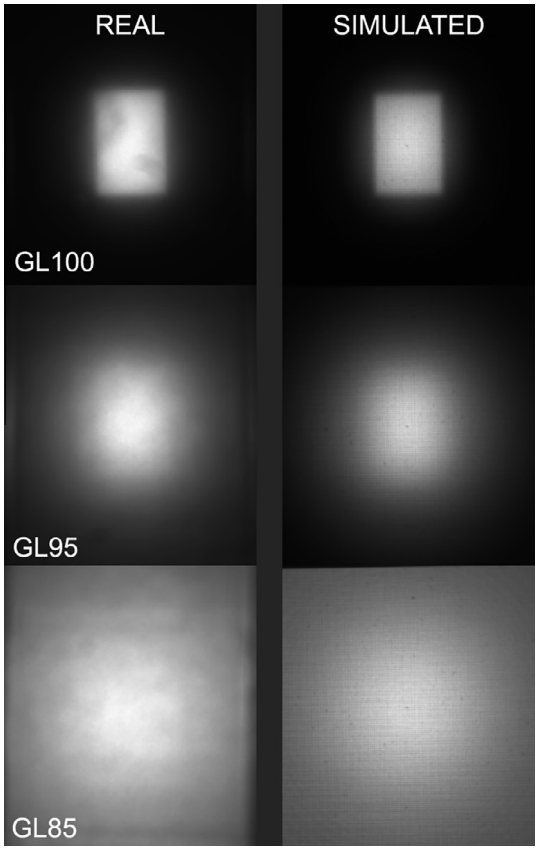

Fig. 16. Images of the real and simulated gloss samples taken with a calibrated digital camera. Note the similarities in the luminance patterns.

6.1. Color perception

Images shown on emissive electronic displays (e.g. LCD, OLED) are widely used as stand-ins for real reflective surfaces in e-commerce, appearance design, and many other applications. Yet looking at an image on a display is not the same as looking at a real reflective surface. For example, Gorzynski and Berns (1990) showed that while observers fully adapt to the chromaticity of the scene illuminant when viewing real colored patches, they only partially adapt when those patches are rendered as images on an emissive display. This incomplete "discounting of the illuminant" for emissive displays can lead to significant errors in color perception that can impact decision-making when emissive displays are used in color-critical applications.

Since one of our goals in designing the surface display is to produce an emissive display has the appearance of a real reflective surface, we replicated Gorzynski and Berns' experiment to investigate if the surface display was seen more as a surface or more as a display. The stimuli used in the experiment are shown in Fig. 17 and consisted of arrays of 21 colored patches seen against a white background. The arrays were presented on the surface display, which was in the multi-source lightbooth. On each trial in the experiment, the observer's task was to choose the most neutral patch in a given array. An adaptive procedure varied the chromaticities of the patches to estimate the observer's "white point", indicating their state of chromatic adaptation.

We studied the effects of two factors (surround illuminant and rendering method) on the observers adaptation state. With respect to surround illuminant, the patch array was always rendered as if it was lit by a tungsten source, but the illumination provided by the lightbooth was set to one of three states: tungsten (consistent);

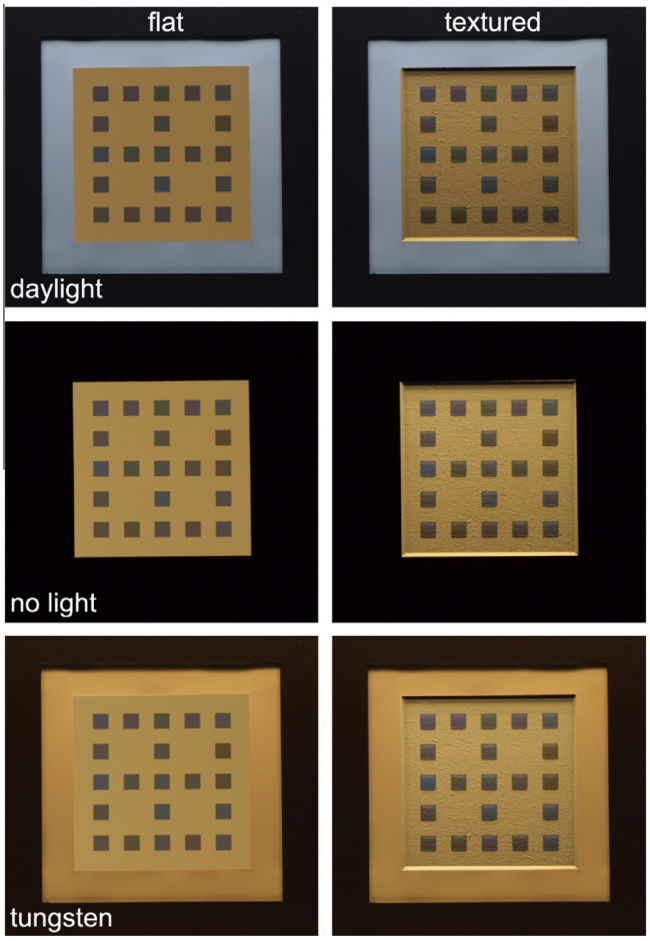

Fig. 17. Stimuli used in the six conditions of our chromatic adaptation experiment. The flat or textured patch arrays were always rendered with respect to a tungsten illuminant. The real-world surround lighting was tungsten, daylight or no light (off). Note that the patch background was white and appeared near-neutral to observers in the no-light and tungsten conditions.

daylight (inconsistent); or no light (neutral). With respect to the rendering method, the patches were always shown on the surface display, but they were either rendered as uniform colored squares (flat), or were rendered using the real-time texturing, shading, and shadowing capabilities of the surface display to make the patches look like real reflective surfaces (textured). Fig. 17 shows the six experimental conditions resulting from the combination of these two factors. Note that, the images in the figure only partially represent what observers saw, and that in particular, in the "no light" and "tungsten" conditions both the immediate backgrounds and surrounds of the patches looked nearly-neutral.

The results of the experiment are shown in Fig. 18a. For comparison, Gorzynski and Berns' results are shown in Fig. 18b. The large triangles in each figure represent the chromaticities of the booth illuminants: daylight (lower left), and tungsten (upper right). The legends under each figure indicate the experimental conditions. To set context, note that in Gorzynski and Berns' experiment, while observers completely adapted to the tungsten booth illuminant when real reflective patches were used (red[1] square), they showed middling levels of adaptation when the patches were rendered on a display. In contrast, our observers showed significantly different adaptation states in the three illumination conditions, however we did not find significant differences between the

[1] For interpretation of color in Fig. 18, the reader is referred to the web version of this article.

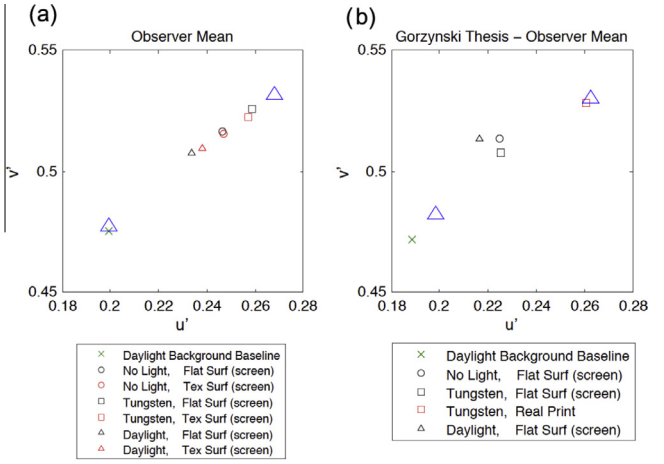

Fig. 18. (a) Results of our chromatic adaptation experiment) compared with (b) the results of Gorzynski and Berns (1990). Note that the results for the surface display show a more systematic pattern of adaptation than the results for the standard emissive display used by Gorzynski and Berns.

"flat" and "textured" rendering conditions. The result indicate that with respect to chromatic adaptation and color perception, the surface display produces results that are significantly more like those found for real surfaces than a standard display.

While this work is preliminary, and we need to do further investigations to understand how the ways in which the surface display appears like a reflective surface (illumination situatedness, responsiveness, realistic materials/lighting) affects the observer's degree chromatic adaptation, this work illustrates how a surface display might facilitate the study color in surface and material perception, by providing an emissive electronic display that is a quantitatively better surrogate for a real surface.

6.2. Gloss perception

Another area where the surface display could contribute to the study of surface and material perception is in the study of gloss. Perceived gloss is a multidimensional property that depends in complex ways on the reflectance and topography of a surface, and the arrangement and structure of light sources in a scene. It also has a dynamic character that is revealed through motion of the surface or movement of the observer. The capabilities of the surface display allow the accurate simulation of all these aspects of glossy surfaces, with the ease of parameter specification and experimental control that come with digital modeling and rendering.

Fig. 19 shows examples of the kinds of glossy stimuli we have generated using the surface display. They are reproductions of seven gray samples from the NCS gloss scale, a physical "book of gloss" whose samples have been measured physically and validated for perceptual uniformity. The images in the figure show of models of the samples rendered to the surface display with accurate geometry and reflectance properties. The samples are also illuminated by a simulated light source that is a geometrically and radiometrically accurate model of the actual lighting in the scene. Because of the real-time capabilities of the system, all the appearance properties of the surfaces change accurately in response to manipulation of the display, movement of the observer, or commands from the experimenter. While we have not yet used the surface display for the study of gloss perception, its capabilities should greatly facilitate the design and execution of display-based psychophysical studies, and because of the realism of the display, those results should transfer well to predict the appearances of real-world surfaces.

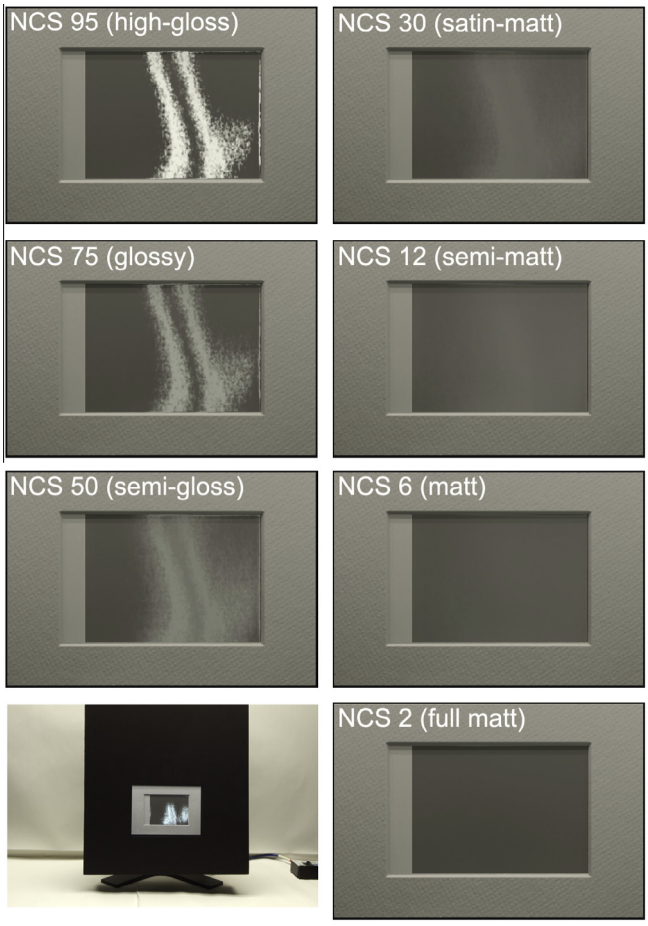

Fig. 19. Simulated gloss samples from the NCS gloss scale. A simulated sample in-situ on the surface display is shown on the lower left.

6.3. Softproofing for digital printing

A concrete example of where these kinds of results would be valuable is in the area of softproofing for digital printing. Users of high-end printing systems or online printing services that offer many choices of inks, papers, and finishing methods often want to preview what a print will look like onscreen before making the actual print. As described in Section 2, a variety of digital softproofing techniques have been developed, but they still often fall short of providing WYSIWYG (what you see is what you get) previewing because of the complex interactions between printing inks, papers, illuminants, and viewing conditions.

We have been developing tools for softproofing using the surface display system. A screenshot from the application is shown in Fig. 20. Here a simulated digital print of a photograph has been rendered to the surface display. The two halves of the simulated print show the effects of different kinds of paper. The left half uses a matte paper with a canvas-like texture, and the right half uses a glossy paper with a smoother texture. While the still image in Fig. 20 illustrates the how the choice of paper will affect the appearance of the final print, the true value of the surface display comes from being able to directly manipulate and interact with the simulated print dynamically under realistic lighting conditions in to understand the effects that media choices will have on the appearance of the final print. While this work is also just beginning, and there is still significant development and validation work to do. This application demonstrates the potential utility of the surface display for not only basic appearance research, but for applied appearance design and engineering as well.

Fig. 20. Composite screenshot from a prototype softproofing application based on the surface display. Note the impact of the matte/canvas and glossy/smooth papers on the appearance of the photograph.

7. Conclusions, limitations, and future work

In this paper we have described our efforts to create a surface display system that produces realistic digital representations of real-world surfaces. Our system supports the accurate simulation of the colors, glosses, and textures of surfaces with complex three-dimensional properties, and allows users to view and interact with these virtual surfaces as if they were real ones.

While we believe this work is significant, there is still much to be done. First, in our current system we model the properties of flat opaque reflective surfaces using the Ward–Dür model. The world is filled with surfaces with more complex geometric and reflectance properties, including effects of mesoscale geometry and translucency, so incorporating more sophisticated surface and light reflection models into our system would be one important effort. Second, in our current system we use a relatively simple lighting environment, and we capture and model it in an offline process. Real-time capture of complex illumination fields would be another worthwhile addition to the system. Next, the user experience could benefit from device-free tracking, and more sophisticated displays that provide stereoscopic and high-dynamic range output, and more tangible form-factors such as tablets and flexible screens. Finally, further validation of the system should be done both physically (particularly with respect to the system's representation of texture) and psychophysically, to determine how well the system conveys the dynamic appearance properties of real surfaces.

Surface display systems such as the one described in this paper represent a powerful and achievable new approach to bridging the real and virtual worlds. One that has important potential applications in science, industry, commerce, education, the arts, or any field where it is important to represent, communicate, and understand the appearance of surfaces. The work described in this paper will hopefully enable further steps in this direction.

Acknowledgment

This work was supported by a Grant from the National Science Foundation (#10644112) to James A. Ferwerda.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.visres.2014.10.016.

References

Bandyopadhyay, D., Raskar, R., & Fuchs, H. (2001). Dynamic shader lamps: Painting on real objects. In *Proc. international symposium on augmented reality (ISAR '01)* (pp. 207–216).

Bimber, O., & Raskar, R. (2005). *Spatial augmented reality*. Wellesley MA: A. K. Peters Ltd..

Burdea, G., & Coiffet, P. (2003). *Virtual reality technology* (2nd ed.). Hoboken, NJ: John Wiley & Sons.

Buxton, W. (2008). Surface and tangible computing, and the small matter of people and design. In *IEEE international solid-state circuits conference digest of technical papers* (Vol. 51, pp. 24–29).

Colbert, M., Premože, S., & François, G. (2010). Importance sampling for production rendering. *ACM SIGGRAPH 2010 Course Notes*.

Cruz-Neira, C., Sandin, D. J., & DeFanti, T. A. (1993). Surround-screen projection-based virtual reality: The design and implementation of the CAVE. *Proceedings ACM SIGGRAPH, 93*, 135–142.

Darling, B. (2013). Ph.D. dissertation, Rochester Institute of Technology. Available through ProQuest/UMI.

Darling, B. A., & Ferwerda, J. A. (2009). The tangiBook: A tangible display system for direct interaction with virtual surfaces. In *Proceedings IS&T 17th color imaging conference* (pp. 260–266).

Darling, B. A., & Ferwerda, J. A. (2012). Seeing virtual objects: Simulating reflective surfaces on emissive displays. In *Proceedings IS-@@-T/SID 20th color imaging conference* (pp. 135–141).

Day, E., Taplin, L., & Berns, R. (2004). Colorimetric characterization of a computer-controlled liquid crystal display. *Color Research & Application, 29*(5), 365–373.

Dür, A. (2006). An improved normalization for the Ward reflectance model. *Journal of Graphics, GPU, and Game Tools, 11*(1), 51–59.

Fitzmaurice, G. W. (1993). Situated information spaces and spatially aware palmtop computers. *Communications of the ACM, 36*(7), 39–49.

Francois, A. R., & Kang, E. Y. (2003). A handheld mirror simulation. In *Proc. int. conf. on multimedia* (pp. 6–9).

Fuchs, M., Raskar, R., Seidel, H. P., & Lensch, H. (2008). Towards passive 6D reflectance field displays. *ACM Transactions on Graphics, 27*(3), 58 (SIGGRAPH 2008).

Gatt, A., Westland, S., & Bala, R. (2006). Increasing the dimensionality of soft proofing: Gloss and material. In *Proceedings of the IS&T 14th color imaging conference* (pp. 292–295).

Gorzynski, M., & Berns, R. (1990). Effects of ambient illumination and image color balance on the perception of neutral in hybrid image display systems. *Proceedings of the SPIE, 1250*, 111–118.

Hall, R. (1989). *Illumination and color in computer generated imagery*. New York: Springer-Verlag.

Hill, B. (2010). Softproofing system for accurate colour matching and study of observer metamerism. In *Proceedings CGIV '10* (pp. 173–180).

Ishii, H., & Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings ACM CHI '97* (pp. 234–240).

Koike, T., & Naemura, T. (2008). BRDF display: Interactive view dependent texture display using integral photography. In *Proceedings ACM IPT/EDT '08*.

Konieczny, J., Meyer, G. W. (2006). Material and color design using projectors. In *Proceedings of the IS&T third European conference on colour in graphics, imaging, and vision* (pp. 438–442).

Konieczny, J., Shimizu, C., Meyer, G. W., & Colucci, D. (2005). A handheld flexible display system. In *Proceedings of IEEE visualization '05* (pp. 75–81).

Laihanen, P. (1994). Exact soft proofing. *Journal of Image Science & Technology, 38*(5), 432–440.

Lanman, D., Wetzstein, G., Hirsch, M., Heidrich, W., & Raskar, R. (2011). Polarization fields: Dynamic light field display using multi-layer LCDs. *Proceedings ACM SIGGRAPH Asia* (SA 2011) (Article No. 186).

Lazzari, M., Mclaughlin, M. L., Jaskowiak, J., Wong, W., & Akbarian, M. (2002). A haptic exhibition of daguerreotype cases for USC's Fisher Gallery. In M. L. McLaughlin, J. Hespanha, & G. Sukhatme (Eds.), *Touch in virtual environments: Haptics and the design of interactive system* (pp. 260–269). Upper Saddle River NJ: Prentice-Hall.

Mahy, M., Van Eycken, L., & Oosterlinck, A. (1994). Evaluation of uniform color spaces developed after the adoption of CIELAB and CIELUV. *Color Research and Application, 19*(2), 105–121.

Masia, B., Wetzstein, G., Didyk, P., & Gutierrez, D. (2013). A survey on computational displays: Pushing the boundaries of optics, computation, and perception. *Computers and Graphics, 37*(8), 1012–1038.

Max, N. L. (1988). Horizon mapping: Shadows for bump-mapped surfaces. *The Visual Computer, 4*(2), 109–117.

Nayar, S. K., Belhumer, P. N., & Boult, T. E. (2004). Lighting sensitive display. *ACM Transactions on Graphics, 23*(4), 963–979 (SIGGRAPH 2004).

Patil, R. A., Fairchild, M. D., & Johnson, G. M. (2004). 3D simulation of prints for improved soft proofing. In *Proceedings of the IS&T 12th color imaging conference* (pp. 193–199).

Peercy, M., Airey, J., & Cabral, B. (1997). Efficient bump mapping hardware. *Proceedings of ACM SIGGRAPH, 97*, 303–306.

Raskar, R., Welch, G., Low, K., & Bandyopadhyay, D. (2001). Shader lamps: Animating real objects with image-based illumination. In *Proc. 12th eurographics workshop on rendering techniques* (pp. 89–102).

Raskar, R., van Baar, J., Beardsley, P., Willwacher, T., Rao, S., & Forlines, C. (2003). iLamps: Geometrically aware and self-configuring projectors. *ACM Transactions on Graphics, 22*(4), 809–818 (SIGGRAPH 2003).

Sloan, P. P., & Cohen, M. F. (2000). Interactive horizon mapping. *Eurographics Workshop on Rendering*, 281–286.

Sutherland, I. E. (1968). A head-mounted three-dimensional display. *Proceedings AIFPS, 33*, 757–764.

Tsang, M., Fitzmaurice, G. W., Kurtenbach, G., Khan, A., & Buxton, W. (2002). Boom chameleon: Simultaneous capture of 3D viewpoint, voice and gesture annotations on a spatially-aware display. *Proceedings ACM UIST, 4*(2), 111–120.

Vickers, D. L. (1974). *Sorcerer's apprentice: Head-mounted display and wand*. Ph.D. dissertation, Univ. of Utah.

Viriyothai, K., & Debevec, P. (2009). Variance minimization light probe sampling. *ACM SIGGRAPH 2009 Poster Proceedings* (Article 92).

Ward, G. J. (1992). Measuring and modeling anisotropic reflection. *Proceedings ACM SIGGRAPH, 92*, 265–272.

Wetzstein, G., Lanman, D., Heidrich, W., & Raskar, R. (2011). Layered 3D: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Transactions on Graphics, 30*(4), 95 (SIGGRAPH 2011).

Yang, R., Huang, X., Li, S., & Jaynes, C. (2008). Toward the light field display: Autostereoscopic rendering via a cluster of projectors. *IEEE Transactions on Visualization and Computer Graphics, 14*(1), 84–96.