

Color appearance in high-dynamic-range imaging

Ahmet Oğuz Akyüz

University of Central Florida
Computer Science Department
Orlando, Florida 32816

and

Max Planck Institute for Biological Cybernetics
Tübingen, Germany 72076

Erik Reinhard

University of Bristol
Department of Computer Science
Bristol BS8 1UB, United Kingdom

Abstract. When viewing images on a monitor, we are adapted to the lighting conditions of our viewing environment as well as the monitor itself, which can be very different from the lighting conditions in which the images were taken. As a result, our perception of these photographs depends directly on the environment in which they are displayed. For high-dynamic-range images, the disconnect in the perception of scene and viewing environments is potentially much larger than in conventional film and photography. To prepare an image for display, luminance compression alone is therefore not sufficient. We propose to augment current tone reproduction operators with the application of color appearance models as an independent preprocessing step to preserve chromatic appearance across scene and display environments. The method is independent of any specific tone reproduction operator and color appearance model (CAM) so that for each application the most suitable tone reproduction operator and CAM can be selected. © 2006 SPIE and IS&T. [DOI: 10.1117/1.2238891]

1 Introduction

The field of high-dynamic-range imaging is rapidly maturing with improved image capture techniques, graphics algorithms' ability to produce arbitrarily large dynamic ranges, and emerging standards in high-dynamic-range file formats.^{1,2,30} However, current monitor technology imposes severe constraints on the range of luminance values that can be displayed. Although high-dynamic-range (HDR) monitors^{4,5} will be more widely available in the near future, currently they are still costly. Therefore, to prepare HDR images for display on conventional display devices, we must bring the range of values to a displayable range, a process called tone reproduction or tone mapping.

As most offices and homes remain equipped with low-dynamic-range display devices, there is an increasing necessity to prepare HDR images for display, for instance, in archival applications.⁶ Tone reproduction algorithms are addressing this requirement by compressing the image to fit to the range of the display device. However, an inherent

problem with dynamic range reduction is that the disconnect between the illumination conditions of the photographed environment and the illumination conditions of the viewing environment is not accounted for. Ignoring this disconnect may render tone-mapped images unnatural and imprecise in terms of their color appearance.


Chromatic adaptation is a mechanism of the human visual system that enables us to adapt to the dominant colors of illumination.⁷ For example, if we are in an environment that has blue illumination, such as sky light at noon, we adapt to this blue color. Therefore, when observing an outdoors image under fluorescent office lighting, the image will have a bluish appearance unless chromatic adaptation effects are considered. However, most chromatic adaptation models account only for the changes in the color of adapting illuminants (i.e., sun and fluorescent light), ignoring other factors that affect adaptation such as the absolute luminance level of the environment, the level of ambient illumination, and the presence of light sources.

Color appearance models (CAMs) describe the environment with a small number of key parameters to predict how colors will appear to the observer. These models consist of two steps. The first step is a chromatic adaptation transform, which accounts for the state of chromatic adaptation of the observer in a given environment. The second step predicts relative and absolute color appearance attributes such as lightness, chroma, hue, brightness, and colorfulness.⁷ Reproducing colors for a particular viewing environment can be achieved by running the full-color appearance model in reverse, while inserting parameters that describe the viewing environment. Color appearance models are significant improvements over chromatic adaptation models, since they incorporate a more complete description of the environment, and thus enable a more precise reproduction of color.

In this paper, we show how to adapt color appearance models such that they may act as preprocessors to any current tone reproduction operator. The only constraint on tone reproduction operators is that the ratios between the three color channels must be preserved. Many current tone repro-

Paper 04062RR received Jun. 7, 2004; revised manuscript received Nov. 7, 2005; accepted for publication Jan. 19, 2006; published online Sep. 15, 2006.

1017-9909/2006/15(3)/033001/12/\$22.00 © 2006 SPIE and IS&T.



Without preservation of color appearance



Color appearance is preserved

Fig. 1 Bottom image is preprocessed using our approach, and is optimized for viewing under fluorescent light on an *sRGB* monitor. Both images had their luminance range reduced with identical parameters using photographic tone-mapping (color online only).¹²

duction operators satisfy this constraint.^{8–12} The color appearance model then only adjusts the image's chromatic channels, while the subsequent tone reproduction operator will compress the luminance channel. We believe that this method is attractive because it enables us to mix-and-match CAMs with tone reproduction operators to suit any specific task.

Figure 1 depicts a sample result, with and without appearance modeling prior to tone mapping. The top image is created without using a CAM and it has a blue color cast. The blue color cast of the bottom image is correctly removed by applying the CIECAM02 model. The bottom image is prepared to be viewed on an *sRGB* monitor under fluorescent lighting.

In the following section, we give an overview of color appearance phenomena, and argue its importance to HDR imaging. Section 3, introduces CAMs and discusses their applicability to our approach. Section 4 provides a brief overview of tone reproduction techniques. We show how to prepare HDR data to make it compliant with color appearance modeling in Sec. 5, and the application of the models in Sec. 6. Section 7 shows results. We conclude our paper with an overall discussion in Sec. 8.

2 Color Appearance Phenomena

The human visual system (HVS) is a highly sensitive and adaptable device for registering and interpreting images. However, it does not operate as a linear light meter, but due to various chemical, electrical, and physical principles it has evolved to employ multiple adaptation mechanisms.¹³

For a linear-light-measuring device one may expect that a change in illumination intensity would yield a similar change in the output of the device. To some extent this is true for the HVS also, especially if this change is small. However, for large differences in illumination, the HVS will also register chromatic changes—even if they are not present in the input. Several color appearance phenomena are known to exist, such as simultaneous contrast; crispening; spreading; the Bezold-Brücke hue shift; and the Abney, Hunt, Helson-Judd, Stevens, Helmholtz-Kohlrausch, and Bartleson-Breneman effects. An exhaustive description of these phenomena is beyond the scope of this paper, but we note that an excellent overview is given by Fairchild.⁷ In the remainder of this section, we highlight some of them to illustrate the nature of color appearance. We focus on effects that are modeled by current CAMs, and are therefore relevant to our work.

Color appearance phenomena are caused by the specific ways in which the human visual system operates. For instance, the L, M, and S cones in the retina exhibit a peak response roughly to the red, green, and blue wavelengths of spectrum. These responses are converted to opponent signals by the neurons of the retina before being transmitted to the brain. According to the opponent colors theory proposed by Hering,¹⁴ light-dark, red-green, and blue-yellow are opponent pairs. Conversion of colors into their opponent pairs induces significant perceptual effects in our visual system.⁷

The luminance of the environment surrounding us has significant impact on the appearance of colors. For instance, a photograph will appear significantly more colorful and vivid if it is viewed under bright sunlight compared to viewing the same photograph indoors. This effect, called the Hunt effect, shows that colorfulness increases with adapting luminance.¹⁵

The brightness of colors also affects their saturation and colorfulness. That is, light and dark colors appear less saturated than average colors with the same chromaticity.¹⁶ This can be explained through the nonlinear response of our photoreceptors to linear changes in light reaching the retina. In particular, a hyperbolic function is often used to model photoreceptor output.¹⁷ As this function is S-shaped, more compression to the signal is applied to small and large intensities than to intermediate intensities. Moreover, this hyperbolic function is independently applicable to the L, M, and S cones. For a series of colors that differ only in luminance but have the same chromatic content, the combined result of the three different cone types will yield a less-saturated impression of this color for high and low luminance levels than for intermediate luminance levels. This effect will be observable for HDR images where the luminance levels of the objects in the captured scene can be arbitrarily high.

With HDR imaging it is possible to photograph very bright or very dim objects without causing over- or under-exposed regions in the final HDR image. Therefore, the

Table 1 Overview of appearance correlates.

Brightness (Q)	The perceived quantity of light emanating from a stimulus.
Lightness (J)	The brightness of a stimulus relative to the brightness of a stimulus that appears white under similar viewing situations.
Colorfulness (M)	The perceived quantity of hue content (difference from gray) in a stimulus. Colorfulness increases with luminance.
Chroma (C)	The colorfulness of a stimulus relative to the brightness of a stimulus that appears white under similar viewing conditions.
Saturation (s)	The colorfulness of a stimulus relative to its own brightness.
Hue (h and e)	The degree to which a stimulus can be described as similar to or different from stimuli that are described as red, green, blue, and yellow.

difference in illumination of the original scene and the viewing environment can potentially be much larger than in conventional digital imaging. Hence, the already mentioned perceptual effects play a more significant role. We, therefore, argue that color appearance should be accounted for in applications that require the display of HDR data on low-dynamic-range (LDR) devices.

3 Color Appearance Models

To predict the appearance of a colored patch, color appearance models normally require the tristimulus values describing the color of the patch as well as parameters to characterize the environment. The latter may, for instance, consist of the relative background luminance, i.e., the ratio of the background luminance to the luminance of the white point, the adapting field luminance, the white point of the scene, the degree of adaptation, and the ambient illumination of the surround.

These inputs are then used to compute appearance correlates, as listed in Table 1. These correlates characterize the perception of the tristimulus values. Hence CAMs provide a quantitative description of how a given color will appear in its environment. If the environment is altered, the same tristimulus values may thus be perceived differently.

CAMs can be grouped into two classes, namely, those that model spatial interactions between neighboring patches or pixels and those that do not. The former category is arguably capable of modeling more complex phenomena of color vision. Two examples of spatial color appearance models are the multiscale observer model¹⁸ and iCAM¹⁹, which also perform dynamic range reduction. Due to their spatially variant nature, these models may show haloing and ringing artifacts,²⁰ a feature also observed in spatially variant tone reproduction operators.²¹ In this paper, we do not prefer the use of spatially variant color appearance

models because their interplay with spatially variant tone reproduction operators is not well understood and considered an issue for further research.

The class of nonspatial CAMs includes Hunt's model,¹⁶ RLAB⁷, CIECAM97s²², ZLAB²³, and CIECAM02²⁴. Many of these models are logical extensions and further developments of each other with CIECAM02 being the most recent incarnation. This model is relatively simple, invertible, and nonspatial and is considered³ an improvement over the CIECAM97 family of CAMs. For these reasons, we adopted the CIECAM02 model for our experimentations, although we also show results obtained with other CAMs. A summary of this model is provided in the appendix. For a more detailed description, and implementation related issues we refer to Moroney *et al.*,²⁴ Li *et al.*,³ Fairchild,⁷ and Hunt.¹⁶

4 Color in Tone Reproduction

Many tone reproduction operators extract and compress luminance values only and leave chromaticities unaffected, which is achieved by keeping the color ratios between color channels constant.²⁵

Preserving chromatic content under luminance compression can be achieved in several different ways. For instance, the input may be converted to a color space that features a separate luminance channel. Range compression may then be applied to this channel only while leaving the chromatic channels unaffected. For instance the Y_{xy} color space is suitable and was used by Reinhard *et al.*¹² A second approach²⁶ is to compute luminance L for each pixel as a linear combination of the red, green, and blue channels: $L=0.2126R+0.7152G+0.0722B$. The resulting L channel is then compressed with some tone reproduction operator, yielding compressed luminances L' . From these data, color image can be reconstructed by preserving the ratio between the red, green, and blue channels (shown here for the red channel):

$$R' = \frac{L'R}{L}$$

Further, the amount of saturation may be controlled in an *ad hoc* manner by including²⁷ an exponent s :

$$R' = L' \left(\frac{R}{L} \right)^s$$

This exponent should be set to 1 if the tone reproduction operator is to be used in combination with a CAM.

5 Image Preparation

To successfully separate color appearance issues from luminance compression, we propose to adjust for color appearance prior to tone reproduction. First, the environment parameters are estimated, then the CAM of choice is applied, and then the inverse model is invoked with viewing environment parameters. Because these steps are likely to alter the dynamic range of the image, we reset the luminance values prior to tone reproduction. An outline of our algorithm is given in Fig. 2. The implications of the luminance reset are discussed in Sec. 8. The result is an image

with chromatic content commensurate with the display environment, but with a retained high dynamic range.

CAMs tend to be calibrated in SI units. In particular, the input must be specified in candela per square meter (cd/m^2). If a specific image is given in the right units, then any of the models can be directly applied to this image.

However, in most practical cases, the image is given in arbitrary units. It is therefore necessary to transform the input from arbitrary units to SI units. In the following, we discuss such transforms for both calibrated and uncalibrated images.

5.1 Calibrated Input Data

A calibrated HDR image can be created from a sequence of exposures taken with a digital camera by changing the exposure time at each image. In this work, we used raw exposures because they are linear and they are not modified by the white balancing of the camera; both of which are desirable properties to obtain an accurate calibration. They are stored in a camera specific color space, which can be converted to XYZ tristimulus space by using the metadata stored with the images.²⁸ From such a sequence an HDR image can be created by dividing each image with its exposure time and summing up all images afterward.^{29,30}

Although these operations result in a linear HDR image, its values are still in relative units. The luminance of a uniform patch in the scene must be known to find the scale factor that will map the HDR image to absolute units. The scale factor can be calculated by placing a calibrated gray card in the scene, measuring its luminance with a photometer, and taking the ratio of the measured luminance with the luminance recorded in the image. By multiplying all pixels with this scale factor, the HDR image can be converted to absolute units, suitable to serve as input to CAMs.

By measuring the luminance and the chromaticity of the gray card, one can also acquire the environment parameters such as the adapting luminance and the reference white. These parameters are also used in the forward color appearance model.

5.2 Uncalibrated Input Data

Estimation of absolute luminances of an uncalibrated HDR image is composed of two steps. First the image is converted to XYZ tristimulus space, and second a scaling factor is found such that when the image pixels are multiplied with it, they approximate real world luminances. The ordering of these steps is not important since HDR images are linear.* Obviously estimating absolute luminances of an uncalibrated input data is an underconstrained problem that cannot be solved without additional information or assumptions. However, in many cases it is possible to infer an approximate range of values in candelas per square meter.

We first convert the linear RGB values of the HDR image to XYZ color space using the $sRGB$ conversion matrix. Our motivation is that most modern digital cameras use $sRGB$ output standard. The forward and inverse conversion matrices are:




Fig. 2 High-level visualization depicting the main steps of our algorithm.

$$M_{sRGB \rightarrow XYZ} = \begin{bmatrix} 0.4124 & 0.3576 & 0.1805 \\ 0.2126 & 0.7152 & 0.0722 \\ 0.0193 & 0.1192 & 0.9505 \end{bmatrix},$$

$$M_{XYZ \rightarrow sRGB} = \begin{bmatrix} 3.2406 & -1.5372 & -0.4986 \\ -0.9689 & 1.8758 & 0.0415 \\ 0.0557 & -0.2040 & 1.0570 \end{bmatrix}.$$

If the original color space of the HDR image was not $sRGB$, this conversion could introduce bias. However, as long as the original color space of the image was in RGB space, this conversion approximates the relative tristimulus values.

For many images, the shape of the histogram provides a clue as to its absolute luminance values. We thus estimate the absolute luminance values of an image by computing its key. The key of an image indicates whether the scene was subjectively light, normal, or dark. Since brightness can be crudely modeled with log-luminance^{8,12,31} we view the normalized log-average luminance of an image to be a useful approximation of its key:

$$\log L_{av} = \frac{1}{N} \sum_{x,y} \log[\delta + L(x,y)],$$

$$k = \frac{\log L_{av} - \log L_{min}}{\log L_{max} - \log L_{min}}.$$

Here, N is the number of pixels in the image, L is the relative luminance for pixel (x,y) , and δ is a small offset to avoid the singularity occurring for black pixels. In the second equation, the log-average luminance is normalized by the dynamic range of the image to compute its key, and L_{min} and L_{max} are the minimum and maximum relative luminances of the image. They are computed after excluding a fraction of lightest and darkest pixels from consideration to make the method robust against outliers. In most cases, excluding 5% of the total number pixels both from the dark and light ends is sufficient.

In many cases, the key k correlates well with the overall dark or light appearance of an image. Hence, we can use k to estimate the overall luminance and compute a scale factor f by which we multiply each pixel. Because k is computed in the log domain, a reasonable approach is to linearly relate f to k :

$$f = 10^4 k / L_{max}.$$

*A properly created HDR image is linear, since the nonlinearities of the camera are inverted during HDR recovery using the inverse of the camera response curve.

Table 2 The input parameters of the CIECAM02 model.

XYZ	Absolute tristimulus values of the stimulus
$X_w Y_w Z_w$	Relative tristimulus values of the white point
L_A	Adapting field luminance in candelas per square meter
Y_b	Relative luminance of the background
Surround	Relative luminance level of the surround specified as dim, dark, or average

The constant 10^4 cd/m² is based on typical maximum luminance values found for different types of scenes.³² By multiplying the input with f , we approximate the real-world luminance values of an uncalibrated input image.

6 Algorithm

In this section, we show the application of the CIECAM02 model. Since the input parameters of RLAB and CIECAM97s are slightly different, the other models may be used with minor modifications.

The CIECAM02 model requires five parameters, shown in Table 2 (see also the appendix). These parameters can be given as user parameters, or may be estimated from the data. The background is the immediate surrounding area of the stimulus. In an image, this corresponds to a neighborhood around each pixel. The adapting field is considered to be outside of the background in the visual field. The surround is defined to be the entire area outside of the background. In practice, it can be considered to be the entire room or the environment in which an image (or scene) is observed.

We set the adapting field luminance to the 20% of the luminance of the white point.^{7,16,33} This assumes⁷ that scenes integrate to a gray with a reflectance factor of 0.2. Y_b is the ratio of the background luminance to the luminance of the white point. For a stimulus surrounded by a uniform background it may be possible to measure Y_b directly. However, in imaging applications the background of every pixel is composed of surrounding pixels, which are likely to change from pixel to pixel. This might require recalculation of the background for every pixel and is deemed to be impractical.⁷ Therefore, in imaging applications the relative background luminance is typically set to 20% of the luminance of the white point by making a gray world assumption.^{7,33}

The relative tristimulus values of the white point may be computed in various ways. For our calibrated images we use the XYZ values measured on the gray card. For many other images we have found that the gray-world assumption holds, which means, any systematic deviation from neutral gray is due to the color of the illuminant. Hence, the image's white point may be estimated from the average X , Y , and Z values. However, for some images the dominant hue

Table 3 Relative white points (XYZ) for common scene types and a selection of CIE illuminants and their associated correlated color temperatures T ; the values in this table are derived from data in Ref. 34.

Scene	T (in K)	X_w	Y_w	Z_w
candle flame	1850	132.614	100.000	11.511
sunrise/sunset	2000	127.432	100.000	14.517
100 W incandescent	2865	109.840	100.000	35.558
tungsten (TV/film)	3200	105.975	100.000	45.347
summer sunlight at noon	5400	97.584	100.000	94.252
summer sun + sky	6504	95.047	100.000	108.883
CIE A (incandescent)	2854	109.840	100.000	35.558
CIE B (direct sunlight)	4874	109.215	100.000	75.199
CIE C (indirect sunlight)	6774	98.071	100.000	118.185
CIE D50 (noon skylight)	5000	96.396	100.000	82.414
CIE D65 (average daylight)	6504	95.047	100.000	108.883
CIE E (normalized reference)	5500	100.000	100.000	100.000
CIE F2 (office fluorescent)	4150	99.187	100.000	67.395

shift is not due to the illuminant, in which case the gray world assumption is violated. In such cases, we use Table 3 to estimate an appropriate white point.

For an outside scene or a normally lit office using an *sRGB* specified monitor the surround can be set to average.³³ However if the lights of the room are turned off so that the illumination is only due to the monitor, the surround can be set to dim or dark.

Based on these settings, the CIECAM02 model specifies a set of parameters, including F , the factor for degree of adaptation (see the appendix). The degree of adaptation D is then determined as a function of both F and L_A . The degree of adaptation is crucial because the adaptation level changes with the luminance level of the surround. Also the degree of chromatic adaptation is less for self-emitting objects. To account for this, we set D to 0 for the lightest pixels and smoothly interpolate it for darker pixels. This improves the appearance of light sources, which are usually the lightest pixels in an HDR image. Note that varying D per pixel does not create a spatial model because D is adjusted according to each pixel's luminance value, not according to each pixel's neighborhood.

We base our estimate for light source detection (and for objects with high luminance) on the key k of the image. We compute a threshold L_T , above which pixels will be assumed as possible light sources:




Fig. 3 The degree of adaptation as the value of s changes.

$$L_T = L_{\min} + [0.6 + 0.4(1 - k)](L_{\max} - L_{\min}).$$

This approximation serves to detect the objects whose appearance will be preserved independent of the chromatic adaptation of the observer. It will typically detect light sources, as well as strong reflections and other image areas that are likely to drive an observer's adaptation.

We use Hermite interpolation to smoothly vary D between 10% below and 10% above this threshold, which corresponds to a range of luminance values between $L_T \pm 0.1(L_{\max} - L_{\min})$. This range is clamped if either boundary exceeds the luminance range of the image:

$$L_{T_0} = \max [L_{\min}, L_T - 0.1(L_{\max} - L_{\min})],$$

$$L_{T_1} = \min [L_{\max}, L_T + 0.1(L_{\max} - L_{\min})].$$

If the luminance of a pixel is below L_{T_0} we do not change the D value computed by the color appearance model. If it is greater than L_{T_1} we set D to 0 for that pixel. For a pixel with luminance L , where $L_{T_0} \leq L \leq L_{T_1}$, the degree of adaptation D is then computed with the aid of interpolation value s :

$$s = \frac{L - L_{T_0} + 0.1(L_{\max} - L_{\min})}{L_{T_1} - L_{T_0}},$$

$$D' = D(1 - 3s^2 + 2s^3).$$

Note that by using Hermite interpolation, as shown in Fig. 3, we gradually drop the degree of adaptation from D to 0. Since we detect light sources according to their luminance, some low-intensity sources will not be detected as light sources. This is appropriate treatment, as these sources do not affect human visual adaptation.

The effect of varying D' is demonstrated in Fig. 4. The left image is created without applying a CAM and therefore shows a yellow color cast in the background. The middle image is obtained by using the CIECAM02 model with chromatic adaptation applied to all pixels. Note that the color cast in the background is removed but the appearance of the light source is also changed. Finally, the interpolation




Fig. 4 Discounting the light source (c) correctly preserves its yellow appearance. Ignoring light source detection (b) overcorrects compared with the uncorrected image (a) (color online only).

scheme is shown in the right image, which depicts an appropriately corrected background as well as a properly discounted light source.

With the XYZ tristimulus values and all environment parameters now available, appearance correlates are computed using the CIECAM02 model. These are then transformed back into XYZ tristimulus space by applying the inverse CIECAM02 model using viewing environment's parameters. In our experiments, we used an *sRGB* monitor in a room lit by fluorescent light to view the images. In this type of mixed chromatic adaptation environment (i.e., light of the monitor, and the fluorescent light), adaptation shifts occur from the monitor's white to the color of the ambient light. In a series of experiments performed by Brainard *et al.*,³⁵ Choh *et al.*,³⁶ and Katoh *et al.*,³⁷ the amount of this shift was found to be as small as 10 to 20% if the observers' eyes were fixated at the monitor, and as large as 40 to 50% if the observers were allowed to let their eyes wander off the monitor. However, even in the latter case when the eye focuses back on the monitor, it quickly restores its adaption back to the monitor's white point.³⁷ For this reason, we primarily used the white point of an *sRGB* monitor as the reference white of the inverse model, unless otherwise indicated.

We set the adapting luminance $L_A = 16 \text{ cd/m}^{-2}$, since the images are prepared to be viewed in a normally lit office.⁷ We set Y_b to 20 since the background of any pixel is not uniform.³³ The surround parameter is set to average since the ratio of the ambient light to the monitor white is greater than 20% in a normally lit office. However, in certain conditions such as viewing a monitor at home, the surround may be considered as dim.⁷

The inverse model yields XYZ tristimulus values that could in principle be converted back to RGB values. Note, however, that one of the steps in the forward CIECAM02 model involves range compression using a hyperbolic (sigmoidal) curve. With a different adapting field luminance L_A the range expansion in the reverse model is incomplete. Hence the dynamic range on output is smaller than on input. Since we aim to separate range compression from color appearance, we therefore convert the output of the inverse model to Y_{xy} color space, and reset the Y values for each pixel to the Y values read on input. This approach affects only the luminance of each pixel but retains the chromaticities computed by the color appearance model. If luminance values are not reset to their original values after chromatic correction, reduction of the dynamic range will be shared between both the color appearance model and the tone re-




Fig. 5 Separate exposures used to create the calibrated HDR image. Each exposure is 1 *f*-stop apart, ranging from 1/4000 to 1/8 s. Although these images are converted to *sRGB* space for display purposes, their raw originals are used to recover the HDR image.

production operator, potentially causing uncontrollable results. The luminance reset is the final step before tone reproduction can be applied.

In our algorithm, we ensure that chromatic effects and range compression are cleanly separated into different algorithms. The CIECAM02 model is now responsible for applying chromatic adaptation and color appearance related effects, whereas any subsequent tone reproduction operator will be responsible for reducing the dynamic range of the image and preserve all luminance-related appearance effects.

7 Results

We performed our experiments on one calibrated HDR image and several uncalibrated HDR images. The calibrated image was created from a sequence of raw exposures captured by a Nikon Coolpix 5400 camera (Fig. 5). An 18% gray card and a SpectraScan PR650 spectrometer were used to measure the adapting luminance, white point, and the color temperature of the photographed scene (Table 4).

Table 4 Environment parameters of the calibrated HDR image as measured by a SpectraScan PR650 spectrometer aimed at a 18% gray card; although color temperature is not directly used by CIECAM02 it is given for illustration.

L_A	240 cd/m ⁻²
$X_w Y_w Z_w$	95.45, 100.00, 129.22
Surround	average
$T(K)$	7969 K




Fig. 6 (a) Image is only tone-mapped with the photographic tone mapping operator.¹² We applied our algorithm on the other images prior to tone mapping with different display white points. The display white point of (b) was set to D65, and the display white points of the lower images were set to a weighted average of D65 and F2 (c) and D65 and A (d). The surround for all three images was set to average.

Our results for the calibrated HDR image are shown in Fig. 6. The image in the upper-left corner is tone-mapped by the photographic tone mapping operator¹² without application of our algorithm. We applied our algorithm to the other three images with the same scene but different display parameters prior to tone mapping. The scene and display parameters are listed in Tables 4 and 5, respectively. For the image in the upper-right corner we set the display white point to D65 and the surround to average, thus assuming an *sRGB* monitor under neutral lighting. The display white for the image in the lower-left corner was set to a weighted average of D65 (60%) and fluorescent light F2 (40%) to represent a mixed state of adaptation between the monitor white and fluorescent office lighting. The surround was set to average. The lower-right image is similar, except now incandescent illumination is assumed. Therefore its white point was set to 60% D65 and 40% incandescent A. The surround parameter was set to average.

In comparison the blue color cast in the upper-left image does not exist in the other three images. Also as the display environment changes from neutral to fluorescent and incandescent, the colors in the image shift toward the color of the white point. With a simple white balance applied on the

Table 5 The viewing environment parameters used for producing images in Fig. 6.

	Upper Right	Lower Left	Lower Right
L_A	16 cd/m ⁻²	16 cd/m ⁻²	16 cd/m ⁻²
$X_w Y_w Z_w$	95.05, 100.00, 108.88	96.71, 100.00, 92.29	100.97, 100.00, 79.55
Surround	average	average	average




Fig. 7 Effect of changing the surround of the display environment. The surround is set to (a) average, (b) dim, and (c) dark. We can observe that as the relative luminance of the surround decreases, the contrast and the colorfulness of the images are increased by the color appearance model. This is performed to counter the actual decrease in contrast and colorfulness.

camera this effect cannot be achieved. White balancing models the chromatic adaptation that takes place in the original scene of the image, but it does not consider the adaptation of the observer in the display environment. Also white balancing in the camera does not distinguish between light sources and reflective objects to adjust the degree of chromatic adaptation accordingly.

The effect of choosing different values for the surround parameter is demonstrated in Fig. 7. The display white for all the images was set to D65, and the surround is set to average, dim, and dark from left to right. The contrast and the colorfulness of the images are increased by the color appearance model with the darkening surround. This is explained by the Hunt effect and Bartleson-Breneman equations, which suggest that colorfulness and image contrast decrease with a decrease in illumination level. CIECAM02 model counters this change by increasing colorfulness and contrast.⁷

In Fig. 8, we demonstrate the accuracy of our light source detection heuristic. The image on the left is just a tone-mapped version of the HDR image. The light sources detected by our algorithm are marked as white pixels on the right image. Although the sky is correctly detected, some small highlights are also detected as light sources. If desired, these highlights can easily be excluded by using a minimum area threshold.

The overall effect of our approach to separate range compression from color appearance issues for uncalibrated high dynamic range images is illustrated in Fig. 9.

In this figure we compare our results with the results of the iCAM (second image) and multiscale observer model (third image), which aim to accomplish tone reproduction and color appearance preservation simultaneously. To create our result (rightmost image) we chose CIECAM02 as the color appearance model and photographic tone-




Fig. 8 Image in (b) depicts the detected light sources in (a) as white pixels.

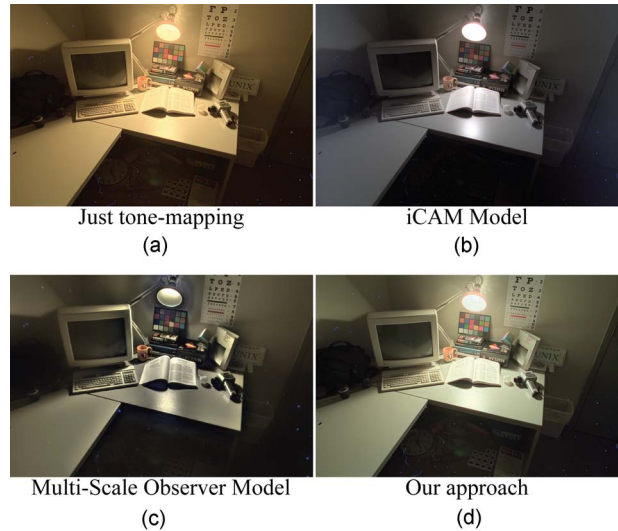


Fig. 9 (a) Tone-mapped image, and (b) output of iCAM model,¹⁹ (c) the multi-scale observer model,¹⁸ and (d) tone-mapped image with appearance preserving preprocessing (our approach).

mapping as the tone reproduction operator.¹² Our result is prepared for viewing under fluorescent light. We estimated the desk lamp in this figure to be close to a CIE A incandescent light source, and therefore chose the corresponding white point (Table 3) as one of the CIECAM02 forward model surround parameters.

For some scenes, dominant hue shift may not be due to the color of the illuminant. In such cases, we estimate the white point using Table 3. For the sunrise image shown in Fig. 10 (right), we estimated the correlated color temperature to be around 3300 K—higher than the table would indicate because half the image contains blue sky, which increases the color temperature. The result looks more plausible than the result produced under the gray world assumption, shown in Fig. 10 (left).

Figure 11 shows results created with and without our approach to color appearance for different tone reproduction operators. In each case, the luminance compression was achieved by applying an operator that does not alter the

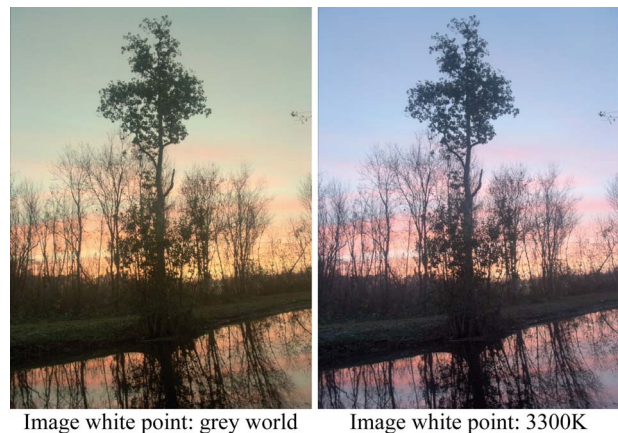


Fig. 10 Comparison of different estimated image white points. The viewing environment is assumed to be lit by fluorescent light.




Fig. 11 Parking garage without (a) and with our preprocessing (b) is compressed with photographic tone reproduction.¹² Ashikhmin's operator,³⁸ bilateral filtering,¹⁰ and histogram adjustment,⁹ from top to bottom. The right column is best viewed under fluorescent light.

ratios between the red, green, and blue channels, such that there is no unwanted interaction between the color appearance model and luminance compression algorithms.

To demonstrate that our algorithm is not bound to a specific color appearance model and tone reproduction operator, we display further results created with different color appearance models and tone reproduction operators in Fig. 12.

Using the CIECAM02 model and the photographic-based tone-mapping in Reinhard *et al.*¹² we show the effectiveness of our approach on a further selection of images in Fig. 13.

8 Discussion

After the application of the forward and backward CAM we reset the luminances of the HDR image to their original values. This step is performed to undo any luminance compression that may be performed by CAMs. For instance, in the forward CIECAM02 model, there is a nonlinear response compression step that is not fully inverted in the backward model unless the same adapting luminance as the forward model is used. Thus, without resetting luminances to their originals, luminance compression will be shared between the CAM and the tone reproduction algorithm, which might yield uncontrollable results. Resetting luminances prior to tone mapping prevents this interaction.




Fig. 12 (a) Tone-mapped image with the photographic tone reproduction operator,¹² (b) the CIECAM97s model applied to HDR data and tone-mapped with the same operator, (c) the tone-mapped image with Tumblin and Rushmeier's operator,⁸ and (d) the RLAB model applied to the high dynamic range data and tone-mapped with the same operator.

Chromatic adaptation models, on the other hand, take into account both the scene and display white points. By using these models, one can compute corresponding colors across two environments that differ only in white point. Simple von Kries type chromatic adaptation models are independent of luminance, which make them unusable to predict color appearance in case of luminance differences.⁷ Some chromatic adaptation models extend von Kries type of adaptation to include additional factors such as adapting luminance and background. Nayatani *et al.* and Fairchild's model are examples of this type.⁷

In HDR imaging, the differences between scene and display environments can potentially be much larger than in conventional digital imaging applications, since there is no limit on the luminance values that can be captured. We therefore argue that using CAMs is essential to preserve appearance of HDR scenes.

In this paper, we show how to preserve color appearance in HDR imaging by combining CAMs with tone reproduction operators. The strengths and weaknesses of our algorithm are the same as its constituent components.

We note that any spatially invariant CAM can be used, as long as it is able to process HDR data. For instance the ZLAB model²³ is restricted to intermediate values of adapting luminance, and is therefore not suitable for HDR data.

Research in tone reproduction is currently very active. New tone reproduction operators become available on a regular basis. We believe that no single operator will be suitable for all possible tasks such as film, photography, computer graphics, and scientific and medical visualization. A more likely scenario is that some operators will be more applicable for certain specific tasks, whereas other operators will be more suitable for other tasks.

These developments make it impractical to target one specific tone reproduction operator for enhancement with techniques borrowed from the field of color appearance.



Fig. 13 Further results with and without appearance preserving pre-processing (for viewing under fluorescent light) (color online only).

Our technique decouples color appearance issues from tone reproduction. Therefore, a key advantage of our technique is that it operates as a preprocess to any subsequent tone reproduction operator as long as the operator compresses the luminance channel only. We believe this method is particularly attractive because HDR imaging will benefit from improvements in both color appearance modeling and tone reproduction.

Appendix CIECAM02

This review of the CIECAM02 color appearance model is based on the technical report³⁹ by CIE. The model requires as input the stimulus, and the parameters of the viewing environment as shown in Table 2.

The stimulus is a uniform patch subtending 2 deg of visual angle and its color is represented by the tristimulus values XYZ . The background is the area adjacent to the stimulus and subtends 10 deg of visual angle. Its relative luminance Y_b is input to the model. The surround is the remainder of the visual field and is normally described as average, dim, or dark. Based on the viewing condition determined by the surround parameter, three parameters may be set according to Table 6. These parameters are c , impact of surround; N_c , chromatic induction factor; and F , factor for degree of adaptation. The input parameter D determines the degree of chromatic adaptation and it is calculated by the following equation:

Table 6 Parameters associated with the viewing condition of the surround.

Viewing Condition	c	N_c	F
Average surround	0.69	1.0	1.0
Dim surround	0.59	0.9	0.9
Dark surround	0.525	0.8	0.8

$$D = F \left[1 - \frac{1}{3.6} \exp \left(\frac{-L_A - 42}{92} \right) \right]$$

The model then proceeds by calculating a number of intermediary parameters that, in turn, will be used for the calculation of appearance correlates.

First, sharpened cone responses RGB are computed from the tristimulus values XYZ through the chromatic adaptation transform (CAT02). This is a linear transformation that achieves von Kries normalization. The computation, the matrix and its inverse are given by:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = M_{\text{CAT02}} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix},$$

$$M_{\text{CAT02}} = \begin{bmatrix} 0.7328 & 0.4286 & -0.1624 \\ -0.7036 & 1.6975 & 0.0061 \\ 0.0030 & 0.0136 & 0.9834 \end{bmatrix},$$

$$M_{\text{CAT02}}^{-1} = \begin{bmatrix} 1.0961 & -0.2789 & 0.1827 \\ 0.4544 & 0.4735 & 0.0721 \\ -0.0096 & -0.0057 & 1.0153 \end{bmatrix}.$$

The sharpened cone responses are then used in the adaptation transform to produce post adaptation cone responses:

$$R_c = [Y_w(D/R_w) + (1 - D)]R,$$

$$G_c = [Y_w(D/G_w) + (1 - D)]G,$$

$$B_c = [Y_w(D/B_w) + (1 - D)]B.$$

A further set of five intermediate parameters are required for the computation of appearance correlates. These parameters are

$$F_L = \text{luminance level adaption factor,}$$

$$n = \text{background induction factor,}$$

$$N_{bb} = \text{brightness background factor,}$$

$$N_{cb} = \text{chromatic background factor,}$$

z = base exponential nonlinearity.

Their computation is given by the following equations:

$$k = 1/(5L_A + 1),$$

$$F_L = 0.2k^4(5L_A) + 0.1(1 - k^4)^2(5L_A)^{1/3},$$

$$n = Y_b/Y_W,$$

$$N_{bb} = 0.725(1/n)^{0.2},$$

$$N_{cb} = N_{bb},$$

$$z = 1.48 + n^{1/2}.$$

The next step is the postadaptation nonlinear response compression. It is performed in Hunt-Pointer-Estevéz space:⁴⁰

$$M_H = \begin{bmatrix} 0.3897 & 0.6890 & -0.0787 \\ -0.2298 & 1.1834 & 0.0464 \\ 0.0000 & 0.0000 & 1.0000 \end{bmatrix},$$

$$M_H^{-1} = \begin{bmatrix} 1.9102 & -1.1121 & 0.2019 \\ 0.3710 & 0.6291 & 0.0000 \\ 0.0000 & -0.0000 & 1.0000 \end{bmatrix}.$$

The computation of compressed cone responses $R'_a G'_a B'_a$ is given by

$$R'_a = \frac{400(F_L R'/100)^{0.42}}{(F_L R'/100)^{0.42} + 27.13} + 0.1,$$

$$G'_a = \frac{400(F_L G'/100)^{0.42}}{(F_L G'/100)^{0.42} + 27.13} + 0.1,$$

$$B'_a = \frac{400(F_L B'/100)^{0.42}}{(F_L B'/100)^{0.42} + 27.13} + 0.1.$$

The compressed cone responses are then used to compute opponent responses:

$$A = [2R'_a + G'_a + (1/20)B'_a - 0.305]N_{bb},$$

$$a = R'_a - 12G'_a/11 + B'_a/11,$$

$$b = (R'_a + G'_a - 2B'_a)/9.$$

This gives us all the intermediary parameters necessary to compute the appearance correlates. The equations that follow then compute the appearance correlates that characterize the XYZ tristimulus values given the specified viewing conditions.

The hue h , e as well as lightness J , brightness Q , chroma C , colorfulness M , and saturation s of the stimulus can be computed by

$$h = \tan^{-1}(b/a),$$

$$e = 0.25 \left[\cos \left(\frac{h\pi}{180} + 2 \right) + 3.8 \right],$$

$$J = 100(A/A_W)^{cz},$$

$$Q = (4/c)(J/100)^{0.5}(A_W + 4)F_L^{0.25},$$

$$t = \frac{(50,000/13)N_c N_{cb} e(a^2 + b^2)^{0.5}}{R'_a + G'_a + (21/20)B'_a},$$

$$C = t^{0.9}(J/100)^{0.5}(1.64 - 0.29^n)^{0.73},$$

$$M = CF_L^{0.25},$$

$$s = 100(M/Q)^{1/2}.$$

For the prediction of color matches across different viewing conditions, lightness-chroma matches are not identical to brightness-colorfulness matches.⁴¹ In practice, lightness and chroma are the appearance correlates of choice. This model can be analytically inverted to compute the XYZ values that would induce the same appearance correlates in a different environment. The details of the inverse model can be found in Hunt¹⁶ and Fairchild.⁷

Acknowledgments

This work is prepared under the funding of i²Lab at University of Central Florida. We thank Greg Ward for making his high dynamic range images available.

References

1. G. W. Larson, "Logluv encoding for full-gamut, high-dynamic range images," *J. Graph. Tools* **3**(1), 15–31 (1998).
2. G. W. Larson and R. A. Shakespeare, *Rendering with Radiance*, Morgan Kaufmann, San Francisco, CA (1998).
3. C. Li, M. R. Luo, R. W. G. Hunt, N. Moroney, M. D. Fairchild, and T. Newman, "The performance of CIECAM02," in *Proc. IS&T/SID 10th Color Imaging Conf.*, pp. 28–32, Scottsdale (Nov. 2002).
4. H. Seetzen, L. A. Whitehead, and G. Ward, "A high dynamic range display using low and high resolution modulators," in *Proc. The Society for Information Display Int. Symp.* (May 2003), pp. 1450–1453.
5. H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs, "High dynamic range display systems," *ACM Trans. Graphics* **23**(3), 760–768 (2004).
6. K. Devlin, A. G. Chalmers, and E. Reinhard, "Displaying digitally archived images," in *Proc. IS&T Archiving Conf.*, The Society of Imaging Science and Technology, (Apr 2004), pp. 157–162.
7. M. D. Fairchild, *Color Appearance Models*, 2nd ed., Wiley, Chicester, UK (2005).
8. J. Tumblin and H. Rushmeier, "Tone reproduction for computer generated images," *IEEE Comput. Graphics Appl.* **13**(6), 42–48 (1993).
9. G. Ward, H. Rushmeier, and C. Piatko, "A visibility matching tone reproduction operator for high dynamic range scenes," *IEEE Trans. Vis. Comput. Graph.* **3**(4), 291–306 (1997).
10. F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high-dynamic-range images," *ACM Trans. Graphics* **21**(3), 257–266 (2002).
11. R. Fattal, D. Lischinski, and M. Werman, "Gradient domain high dynamic range compression," *ACM Trans. Graphics* **21**(3), 249–256 (2002).
12. E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, "Photographic tone reproduction for digital images," *ACM Trans. Graphics* **21**(3), 267–276 (2002).
13. S. E. Palmer, *Vision Science: Photons to Phenomenology*, The MIT Press, Cambridge, MA (1999).

14. E. Hering, *Outlines of a Theory of the Light Sense*, Harvard Univ. Press, Cambridge (1920).
15. R. Hunt, "Light and dark adaptation and the perception of color," *J. Opt. Soc. Am.* **42**, 190–199 (1952).
16. R. Hunt, *The Reproduction of Colour*, 6th ed., Wiley, Chichester, UK (2005).
17. K. I. Naka and W. A. H. Rushton, "S-potentials from luminosity units in the retina of fish (cyprinidae)," *J. Physiol. (London)* **185**, 587–599 (1966).
18. S. N. Pattanaik, J. A. Ferwerda, M. D. Fairchild, and D. P. Greenberg, "A multiscale model of adaptation and spatial vision for realistic image display," in *SIGGRAPH 98 Conf. Proc.*, pp. 287–298 (July 1998).
19. M. D. Fairchild and G. M. Johnson, "Meet iCAM: an image color appearance model," in *Proc. IS&T/SID 10th Color Imaging Conf.*, pp. 33–38, Scottsdale, AZ (2002).
20. N. Moroney and I. Tastl, "A comparison of retinex and iCAM for scene rendering," *J. Electron. Imaging* **13**(1), 139–145 (2004).
21. K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman, "Spatially nonuniform scaling functions for high contrast images," in *Proc. Graphics Interface '93*, Morgan Kaufmann, pp. 245–253 (May 1993).
22. CIE, "The CIE 1997 interim colour appearance model (simple version), CIECAM97s," Technical report, CIE Pub. 131, Vienna (1998).
23. M. D. Fairchild, "The zlab color appearance model for practical image reproduction applications," in *CIE Expert Symposium '97, Colour Standards for Imaging Technology*, CIE Pub. x014, pp. 89–94, Scottsdale (1997).
24. N. Moroney, M. D. Fairchild, R. W. G. Hunt, C. J. Li, M. R. Luo, and T. Newman, "The CIECAM02 color appearance model," in *Proc. IS&T 10th Color Imaging Conf.*, pp. 23–27 Scottsdale, AZ (2002).
25. C. Schlick, "Quantization techniques for the visualization of high dynamic range pictures," in *Photorealistic Rendering Techniques*, P. Shirley, G. Šakas, and S. Müller, Eds., pp. 7–20, Springer-Verlag New York (1994).
26. ITU (International Telecommunication Union, ITU-R Recommendation BT.709, Basic Parameter Values for the HDTV Standard for the Studio and for International Programme Exchange, ITU, Geneva (1990), formerly CCIR Rec. 709).
27. S. N. Pattanaik, J. Tumblin, H. Yee, and D. P. Greenberg, "Time-dependent visual adaptation for fast realistic image display," in *SIGGRAPH'00: Proc. 27th Ann. Conf. on Computer Graphics and Interactive Techniques*, pp. 47–54 (2000).
28. Adobe Systems Incorporated, "Digital negative DNG specification" (2005).
29. P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *SIGGRAPH 97 Conf. Proc.*, Annual Conference Series, pp. 369–378 (Aug. 1997).
30. E. Reinhard, G. Ward, P. Debevec, and S. Pattanaik, *High Dynamic Range Imaging: Acquisition, Display, and Image Based Lighting*, Morgan Kaufmann, San Francisco, CA (2005).
31. G. Ward, "A contrast-based scalefactor for luminance display," in *Graphics Gems IV*, P. Heckbert, Ed., pp. 415–421. Academic Press, Boston (1994).
32. B. A. Wandell, *Foundations of Vision*, Sinauer, Sunderland, MA (1995).
33. N. Moroney, "Usage guidelines for CIECAM97s," in *Proc. Conf. on Image Processing, Image Quality, Image Capture Systems (PICS-00)*, pp. 164–168, Society for Imaging Science and Technology, Springfield, VA (2000).
34. G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd ed., Wiley, New York (2000).
35. D. H. Brainard and K. Ishigami, "Factors influencing the appearance of crt colors," in *Proc. 3rd IS&T/SID Color Imaging Conf.*, pp. 62–66 (1995).
36. H. K. Choh, D. S. Park, C. Y. Kim, and Y. S. Seo, "Effects of ambient illumination on the appearance of CRT colors," in *Proc. 4th IS&T/SID Color Imaging Conf.*, pp. 224–226 (1996).
37. N. Katoh, K. Nakabayashi, M. Ito, and S. Ohno, "Effect of ambient light on color appearance of softcopy images: mixed chromatic adaptation for self-luminous displays," *J. Opt. Imaging* **7**(4), 794–806 (1998).
38. M. Ashikhmin, "A tone mapping algorithm for high contrast images," in *Proc. 13th Eur. Workshop on Rendering*, Eurographics Association, pp. 145–155 (2002).
39. CIE, "A colour appearance model for colour management systems: CIECAM02," Technical report, CIE Pub. 159, Vienna (2004).
40. R. W. G. Hunt and M. R. Pointer, "A colour-appearance transform for the 1931 standard colorimetric observer," *Color Res. Appl.* **10**, 165–179 (1985).
41. Y. Nayatani, K. Takahama, H. Sobagaki, and K. Hashimoto, "Color-appearance model and chromatic adaptation transform," *Color Res. Appl.* **15**, 210–221 (1990).



imaging.

Ahmet Oğuz Akyüz is a PhD candidate with the School of Computer Science, University of Central Florida, and holds a fellowship position at the Max Planck Institute on Biological Cybernetics in Tübingen. He is working under supervision of Erik Reinhard and Heinrich H. Buelthoff. He received his BS degree from the Middle East Technical University in Ankara, Turkey. His current research interests include color, color imaging, and high-dynamic-range



Erik Reinhard is lecturer with the University of Bristol and holds a courtesy appointment as assistant professor at the University of Central Florida. He received his BS degree and a TWAIO diploma in computer science from Delft University of Technology and his PhD in computer science from the University of Bristol. He was a postdoctoral researcher with the University of Utah. He coauthored the first book on high-dynamic-range imaging (Morgan Kaufmann, 2005) and is founder and co-editor-in-chief of the journal *ACM Transactions on Applied Perception* and guest editor of a special issue on parallel graphics and visualization for the journal *Parallel Computing* (March 2003). He is also coeditor of *Practical Parallel Rendering* (A.K. Peters, 2002). His current interests are in high-dynamic-range imaging and visual perception and their application to problems in computer graphics and computer vision.