iCAM framework for image appearance, differences, and quality

Mark D. Fairchild
Garrett M. Johnson
Rochester Institute of Technology
Chester F. Carlson Center for Imaging Science
Munsell Color Science Laboratory
54 Lomb Memorial Drive
Rochester, New York 14623-5604
E-mail: mdf@cis.rit.edu

Abstract. *Traditional color appearance modeling has recently matured to the point that available, internationally recommended models such as CIECAM02 are capable of making a wide range of predictions, to within the observer variability in color matching and color scaling of stimuli, in somewhat simplified viewing conditions. It is proposed that the next significant advances in the field of color appearance modeling and image quality metrics will not come from evolutionary revisions of colorimetric color appearance models alone. Instead, a more revolutionary approach will be required to make appearance and difference predictions for more complex stimuli in a wider array of viewing conditions. Such an approach can be considered image appearance modeling, since it extends the concepts of color appearance modeling to stimuli and viewing environments that are spatially and temporally at the level of complexity of real natural and man-made scenes, and extends traditional image quality metrics into the color appearance domain. Thus, two previously parallel and evolving research areas are combined in a new way as an attempt to instigate a significant advance. We review the concepts of image appearance modeling, present iCAM as one example of such a model, and provide a number of examples of the use of iCAM in image reproduction and image quality evaluation. © 2004 SPIE and IS&T.* [DOI: 10.1117/1.1635368]

1 Introduction

The fundamental theme of this research can be considered image measurement, and the application of those measurements to image rendering and image quality evaluation. Consideration of the history of image measurement helps set the context for the formulation and application of image appearance models, a somewhat natural evolution of color appearance, spatial vision, and temporal vision models when they are considered in a holistic sense, rather than as individual research fields. Early imaging systems were either not scientifically measured at all, or measured with systems designed to specify the variables of the imaging system itself. For example, densitometers were developed for measuring photographic materials with the intent of specifying the amounts of dye or silver produced in the film. In printing, similar measurements would be made for

the inks as well as measures of the dot area coverage for halftone systems. In electronic systems like television, system measurements such as signal voltages were used to colorimetrically quantify the imaging system.[1] It should be noted that vision-based measurements of imaging systems for image quality do have a long history, as illustrated by the example of Schade's pioneering work.[2] As imaging systems evolved in complexity and openness, the need for device-independent image measures became clear.

1.1 Image Colorimetry

Electronic imaging systems, specifically the development of color television, prompted the first application of device-independent color measurements of images. Wright, in fact, points out that color television could not have been invented without colorimetry.[3] Device-independent color measurements are based on the internationally standardized CIE system of colorimetry first developed in 1931. CIE colorimetry specifies a color stimulus with numbers proportional to the stimulation of the human visual system, independent of how the color stimulus was produced. The CIE system was used very successfully in the design and standardization of color television systems (including recent digital television systems).

Application of CIE colorimetry to imaging systems became much more prevalent with the advent of digital imaging systems and, in particular, the use of computer systems to generate and proof content ultimately destined for other media, such as print. As color-capable digital imaging systems (from scanners and cameras, through displays, to various hardcopy output technologies) became commercially available in the last two decades, it was quickly recognized that device-dependent color coordinates (such as monitor RGB and printer CMYK) could not be used to specify and reproduce color images with accuracy and precision. An additional factor was the open-systems nature of digital imaging in which the input, display, and output devices might be produced by different manufacturers, and one source could not control color through the entire process. The use of CIE colorimetry to specify images across various devices promised to solve some of the new color

reproduction problems created by open, digital systems. The flexibility of digital systems also made it possible and practical to perform colorimetric transformations on image data in attempts to match the colors across disparate devices and media.

Research on imaging device calibration and characterization has spanned the range from fundamental color measurement techniques to the specification of a variety of devices including CRT, LCD, and projection displays, scanners and digital cameras, and various film recording and print media. Some of the concepts and results of this research have been summarized by Berns.[4] Such capabilities are a fundamental requirement for research and development in color and image appearance. Research on device characterization and calibration provides a means to tackle more fundamental problems in device-independent color imaging. For example, conceptual research on design and implementation of device-independent color imaging,[5] gamut mapping algorithms to deal with the reproduction of desired colors that fall outside the range that can be obtained with a given imaging device,[6] and computer graphics rendering of high-quality spectral images that significantly improve the potential for accurate color in rendered scenes.[7] This type of research built on, and contributed to, research on the development and testing of color appearance models for cross-media image reproduction.

1.2 Color Difference Equations

Color difference research has culminated with the recently published CIEDE2000 color difference formula.[8] A color difference equation allows for the mapping of physically measured stimuli into perceived differences. At the heart of such color difference equations lies some form of uniform color space. The CIE initially recommended two such color spaces in 1976, CIELAB and CIELUV. Both spaces were initially described as interim color spaces, with the knowledge that they were far from complete. More than 25 years later, these spaces are still the CIE recommendations, although CIELUV has fallen out of favor.

With a truly uniform color space, color differences can then be taken to be a simple measure of distance between two colors in the space, such as CIE ΔE_{ab}^*. The CIE recognized the nonuniformity of the CIELAB color space, and formulated more advanced color difference equations such as CIE DE94 and CIEDE2000. These more complicated equations are very capable of predicting perceived color differences of simple color patches.

1.3 Image Difference

The CIE color difference formulas were developed using simple color patches in controlled viewing conditions. There is no reason to believe that they are adequate for predicting color difference for spatially complex image stimuli. The S-CIELAB model was designed as a spatial preprocessor to the standard CIE color difference equations, to account for complex color stimuli such as halftone patterns.[9] Spatial preprocessing uses separable convolution kernels to approximate the contrast sensitivity functions (CSF) of the human visual system. The CSF serves to remove information that is imperceptible to the visual system. For instance, when viewing halftone dots at a certain distance, the dots tend to blur and integrate into a single

color. A pixel-by-pixel color difference calculation between a continuous image and a halftone image would result in very large errors, while the perceived difference might in fact be small. The spatial preprocessing would blur the halftone image, so that it more closely resembles the continuous tone image.

S-CIELAB represents the first incarnation of an image difference model based on the CIELAB color space and color difference equations. Recently, this model has been refined and extended into a modular framework for image color difference calculations.[10] This framework refines the CSF equations from the S-CIELAB model, and adds modules for spatial frequency adaptation, spatial localization, and local and global contrast detection. This framework is discussed in more detail below.

1.4 Color Appearance

Unfortunately, fundamental CIE colorimetry does not provide a complete solution for image specification. CIE colorimetry is only strictly applicable to situations in which the original and reproduction are viewed in identical conditions. By their very nature, the images produced or captured by various digital systems are examined in widely disparate viewing conditions, from the original captured scene, to a computer display in a dim room, to printed media under a variety of light sources, to projection displays in dark rooms. Thus color appearance models were developed to extend CIE colorimetry to the prediction of color appearance (not just color matches) across changes in media and viewing conditions (not just within a single condition). Color appearance modeling research applied to digital imaging systems was very active throughout the 1990s, culminating with the recommendation of the CIECAM97s model in 1997[11] and its revision, CIECAM02, in 2002.[12] Details on the evolution, formulation, and application of color appearance models can be found in Fairchild.[13] The development of these models was also enabled by visual experiments performed to test the performance of published color-appearance models in realistic image reproduction situations.[14] Such research on color appearance modeling in imaging applications naturally highlighted the areas that are not adequately addressed for spatially complex image appearance and image quality problems.

1.5 Image Appearance and Image Quality

Color appearance models account for many changes in viewing conditions, but are mainly focused on changes in the color of the illumination (white point), the illumination level (luminance), and surround relative luminance. Such models do not directly incorporate any of the spatial or temporal properties of human vision and the perception of images. They essentially treat each pixel of an image (and each frame of a video) as completely independent stimuli. A review of some current work in the area provides context.

Visual adaptation to scenes and images is not only spatially localized according to some low-pass characteristics, but also temporally localized in a similar manner. To predict the appearance of digital video sequences, particularly those of high-dynamic range, the temporal properties of light and chromatic adaptation must be considered. To pre-

dict the quality (or image differences) of video sequences, temporal filtering to remove imperceptible high-frequency temporal modulations (imperceptible flicker) must be added to the spatial filtering that removes imperceptible spatial artifacts (e.g., noise or compression artifacts).

It is easy to illustrate that adaptation has a significant temporal low-pass characteristic. For example, if one suddenly turns on the lights in a darkened room (as when first awakening in the morning), the increased illumination level is at first dazzling to the visual system, essentially overexposing it. After a short period of time, the visual system adapts to the new, higher level of illumination and normal visual perception becomes possible. The same is true when going from high levels of illumination to low levels (imagine driving into a tunnel in the daytime). Fairchild and Reniff[15] and Rinner and Gegenfurtner[16] have made detailed measurements of the time course of chromatic adaptation. These results suggest temporal integration functions that could be used in models of moving image appearance, and also illustrate one of the mechanisms for spatially low-pass adaptation stimuli due to the influence of ever-present eye movements. Such adaptation stimuli are used in the model described in this work.

There has been significant research on video quality and video quality metrics, often aimed at the creation and optimization of encoding/compression/decoding algorithms such as MPEG2 and MPEG4. By analogy, the still-image visible differences predictor of Daly[17] is quite applicable to the prediction of the visibility of artifacts introduced into still images by JPEG image compression. The Daly model was designed to predict the probability of detecting an artifact (i.e., is the artifact above the visual threshold). The iCAM work reviewed here and elsewhere[18,19] has had a different objective with respect to image quality. Instead of focusing on threshold differences in quality, the focus has been on the prediction of image quality scales (e.g., scales of sharpness, contrast, graininess) for images with changes well above threshold. Such suprathreshold image differences are a different domain of image quality research, based on image appearance that separate the iCAM model from previous image quality models.

Likewise, a similar situation exists in the area of video quality metrics. Metrics have been published to examine the probability of detection of artifacts in video (i.e., threshold metrics), but there appears to be no models of video image appearance designed for rendering video and predicting the magnitudes of perceived differences in video sequences. The latter is one of the ultimate goals of the development of iCAM. Two well-known video image quality models, the Sarnoff just noticeable difference (JND) model and the NASA digital video quality (DVQ) model, are briefly described next to contrast their capabilities with the proposed extensions to the iCAM model.

The Sarnoff JND model is the basis of the JNDmetrix software package (see www.jndmetrix.com) and related video quality hardware. The model is briefly described in a technical report published by Sarnoff,[19] and is more fully disclosed in other publications.[20] It is based on the multiscale model of spatial vision published by Lubin,[21,22] with some extensions for color processing and temporal variation. The Lubin model is similar in nature to the Daly model mentioned before, in that it is designed to predict the probability of detection of artifacts in images. These are threshold changes in images often referred to as just noticeable differences, or JNDs. The Sarnoff JND model has no mechanisms of chromatic and luminance adaptation, as are included in the iCAM model. The input to the Sarnoff model must first be normalized, which can be considered a very rudimentary form of adaptation. The temporal aspects of the Sarnoff model are also not aimed at predicting the appearance of video sequences, but rather at predicting the detectability of temporal artifacts. As such, the model only uses two frames (four fields) in its temporal processing. Thus, while it is capable of predicting the perceptibility of relatively high-frequency temporal variation in the video (flicker), it cannot predict the visibility of low-frequency variations that would require an appearance-oriented, rather than JND-oriented, model. The Sarnoff model also is not designed for rendering video. This is not a criticism of the model formulation, but an illustration of how the objective of the Sarnoff JND model is significantly different from that of the iCAM model. While it is well accepted in the vision science literature that JND predictions are not linearly related to suprathreshold appearance differences, it is certainly possible to use a JND model to try to predict suprathreshold image differences, and the Sarnoff JND model has been applied with some success to such data.

A similar model, the DVQ metric has been published by Watson, Hu, and McGowan[23,24] of NASA. The DVQ metric is similar in concept to the Sarnoff JND model, but significantly different in implementation. Its spatial decomposition is based on the coefficients of a discrete cosine transformation (DCT), making it amenable to hardware implementation, and likely making it particularly good at detecting artifacts introduced by DCT-based video compression algorithms. It also has a more robust temporal filter that should be capable of predicting a wider array of temporal artifacts. Like the Sarnoff model, the DVQ metric is aimed at predicting the probability of detection of threshold image differences. The DVQ model also includes no explicit appearance processing through spatial or temporal adaptation, or correlates of appearance attributes, and therefore also cannot be used for video rendering. Again, this is not a shortcoming, but rather a property of the design objectives for the DVQ model.

While color appearance modeling has been successful in facilitating device-independent color imaging and is incorporated into modern color management systems, there remains significant room for improvement and extension of capabilities. To address these issues with respect to spatial properties of vision and image perception (localized adaptation and spatial filtering) and image quality, the concept of image appearance models has been recently introduced and implemented.[25,26] These models combine attributes of color appearance models with attributes of spatial vision models that have been previously used for image quality metrics in an attempt to further extend the capabilities of color-appearance models. Historically, color-appearance models largely ignored spatial vision (e.g., CIECAM97s), while spatial vision models for image quality largely ignored color.[17,21] One notable exception, and the theme of this special issue of the *Journal of Electronic Imaging*, has been the Retinex model[27–30] and its various derivatives.[31–33] The spatial ATD model[34] and the

S-CIELAB model[9] also address some of these issues to various extents. While the Retinex model was never designed as a complete model of image appearance and quality, its spatially variable mechanisms of chromatic adaptation and color constancy serve some of the same purposes in image rendering, and provide some of the critical groundwork for image-appearance modeling.

The goal in developing an image appearance model has been to bring these research areas together to create a single model applicable to image appearance, image rendering, and image quality specifications and evaluations. One such model for still images, referred to as iCAM, has recently been published by Fairchild and Johnson[25] and is detailed in this work. This model was built on previous research in uniform color spaces,[35] the importance of image surround,[36] algorithms for image difference and image quality measurement,[19,37] insights into observers eye movements while performing various visual imaging tasks and adaptation to natural scenes,[38,39] and an earlier model of spatial and color vision applied to color appearance problems and high-dynamic-range (HDR) imaging.[40] The structure of the iCAM model and examples of its implementation for image appearance are presented next.

1.6 Color and Image Appearance Models

A model capable of predicting perceived color difference between complex image stimuli is a useful tool, but has some limitations. Just as a color appearance model is necessary to fully describe the appearance of color stimuli, an image appearance model is necessary to describe spatially complex color stimuli. Color appearance models allow for the description of attributes such as lightness, brightness, colorfulness, chroma, and hue. Image appearance models extend on this to also predict such attributes as sharpness, graininess, contrast, and resolution.

A uniform color space also lies in the heart of an image appearance model. The modular image difference framework allows for great flexibility in the choice of color spaces. Examples are the CIELAB color space, similar to S-CIELAB, the CIECAM02 color appearance model, or the IPT color space.[12,35] Thus the modular image difference framework can be implemented within the iCAM model, as described in this work, to create a full image appearance and image difference model. It could also be implemented in other color spaces if desired. This is one of the main benefits of its modularity.

Models of image appearance can be used to formulate multidimensional models of image quality. For example, it is possible to take weighted sums of various appearance attributes to determine a metric of overall image quality, as described by Keelen[41] and Engledrum.[42] Essentially, these models can augment or replace human observations to weight image attributes with overall appearances of quality. For instance, a model of quality might involve weighted sums of tonal balance, contrast, and sharpness. A first step toward this type of model is illustrated in more detail later.

2 iCAM Framework

Figure 1 presents a flow chart of the general framework for the iCAM image appearance model as applied to still images originally presented by Fairchild and Johnson.[25] Future updates to the model, along with example images and source code can be found at www.cis.rit.edu/mcsl/iCAM. For input, the model requires colorimetrically characterized data for the image (or scene) and surround in absolute luminance units. The image is specified in terms of relative CIE XYZ tristimulus values. The adapting stimulus is a low-pass filtered version of the CIE XYZ image that is also tagged with absolute luminance information necessary to predict the degree of chromatic adaptation. The absolute luminances (Y) of the image data are also used as a second low-pass image to control various luminance-dependant aspects of the model intended to predict the Hunt effect (increase in perceived colorfulness with luminance) and the Stevens effect (increase in perceived image contrast with luminance). Last, a low-pass, luminance (Y) image of significantly greater spatial extent is used to control the prediction of image contrast that is well established to be a function of the relative luminance of the surrounding conditions (Bartleson and Breneman equations). Refer to Fairchild[13] for a full discussion of the various image appearance effects mentioned earlier and detailed specifications of the data required. The specific low-pass filters used for the adapting images depend on viewing distance and application. Additionally, in some image rendering circumstances, it might be desirable to have different low-pass adapting images for luminance and chromatic information, to avoid desaturation of the rendered images due to local chromatic adaptation (decrease in visual sensitivity to the color of the stimulus). This is one example of application dependence. Local chromatic adaptation might be appropriate for image difference or image quality measurements, but inappropriate for image rendering situations.

The first stage of processing in iCAM is to account for chromatic adaptation. The chromatic adaptation transform embedded in the recently published CIECAM02 model[12] has been adopted in iCAM, since it was well researched and established to have excellent performance with all available visual data. It is also a relatively simple chromatic adaptation model, amenable to image processing applications. The chromatic adaptation model, given in Eqs. (1)–(6), is a linear von Kries normalization of RGB image signals to the RGB adaptation signals derived from
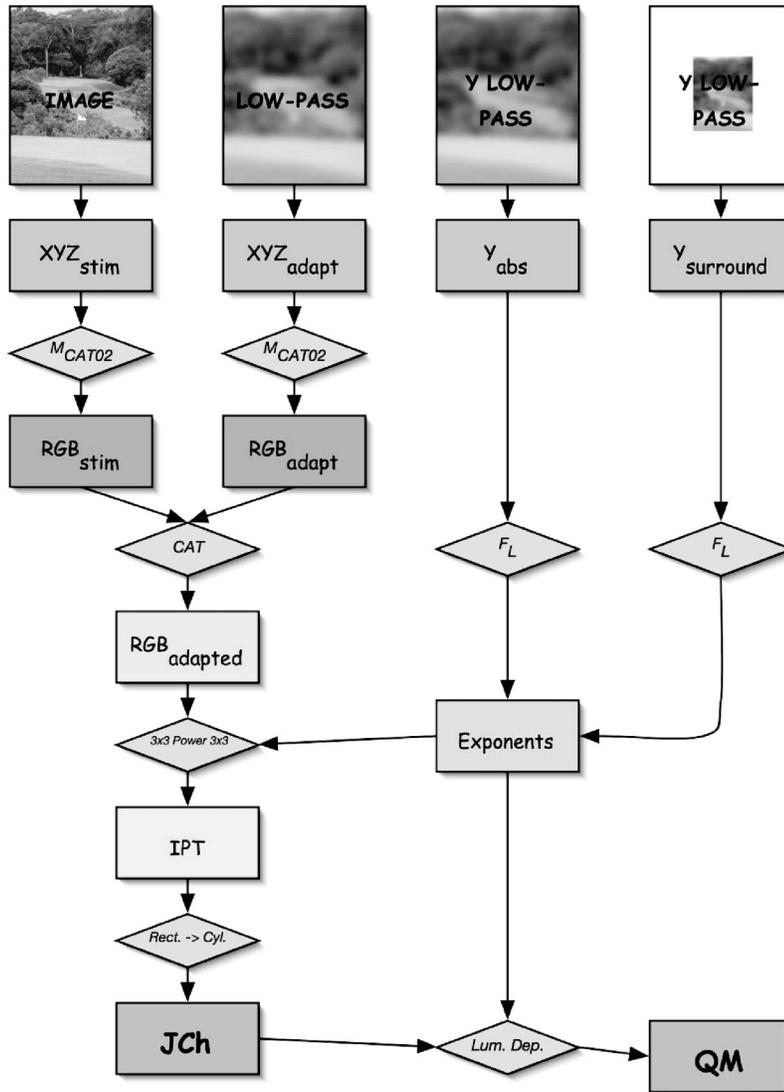
$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \mathbf{M}_{\text{CAT02}} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}, \tag{1}$$

$$\mathbf{M}_{\text{CAT02}} = \begin{bmatrix} 0.7328 & 0.4296 & -0.1624 \\ -0.7036 & 1.6975 & 0.0061 \\ 0.0030 & 0.0136 & 0.9834 \end{bmatrix}, \tag{2}$$

$$D = F \left[1 - \left(\frac{1}{3.6} \right) \exp\left(\frac{-L_A - 42}{92} \right) \right], \tag{3}$$

$$R_c = \left[\left(100 \frac{D}{R_W} \right) + (1 - D) \right] R, \tag{4}$$

$$G_c = \left[\left(100 \frac{D}{G_W} \right) + (1 - D) \right] G, \tag{5}$$

Fig. 1 Flowchart of the iCAM image appearance model. Inputs to the model are CIE tristimulus values XYZ for the stimulus image or scene and a low-pass version used as an adapting stimulus and absolute luminance information for the low-pass image and surround. Adapted signals are computed using the linear chromatic adaptation transform from CIECAM02, and are then converted into an opponend space IPT using the luminance information to modulate a compressive nonlinearity. The rectangular IPT coordinates are then converted to cylindrical correlates of lightness *J*, chroma *C*, and hue *h*. The lightness and chroma correlates can then be scaled by a function of the absolute luminance information to provide correlates of brightness *Q* and colorfulness *M*. If desired, a saturation correlate can be computed as the ratio of chroma to lightness (or colorfulness to brightness).

$$B_c = \left[\left(100 \frac{D}{B_W} \right) + (1-D) \right] B, \qquad (6)$$

with the low-pass adaptation image at each pixel location $(R_W G_W B_W)$. The RGB signals are computed using a linear transformation from XYZ to RGB, derived by CIE TC8-01 in the formulation of CIECAM02. This matrix transformation has come to be called the $\mathbf{M}_{\text{CAT02}}$ matrix, where CAT stands for chromatic adaptation transform. The von Kries normalization is further modulated with a degree-of-adaptation factor *D* that can vary from 0.0 for no adaptation to 1.0 for complete chromatic adaptation. Equation (3) is provided in the CIECAM02 formulation, and is used in

iCAM for computation of *D* as a function of adapting luminance L_A for various viewing conditions. Alternatively, the *D* factor can be established manually. The chromatic adaptation model is used to compute corresponding colors for CIE Illuminant D65, which are then used in the later stages of the iCAM model. This is accomplished by taking the adapted signals for the viewing condition $R_C G_C B_C$ and then inverting Eqs. (1)–(6) for an illuminant D65 adapting white point, and with $D=1.0$. It should be noted that, while the adaptation transformation is identical to that in CIECAM02, the iCAM model is already significantly different, since it uses spatially modulated image data as input rather than single color stimuli and adaptation points. It

also differs completely in the remainder of the formulation, although using CIECAM02 equations where appropriate. One example of this is the modulation of the absolute luminance image and surround luminance image using the F_L function from CIECAM02 given in Eq. (7). This function, slowly varying with luminance, has been

$$F_L = 0.2 \left[\frac{1}{(5L_A + 1)} \right]^4 (5L_A)$$
$$+ 0.1 \left\{ 1 - \left[\frac{1}{(5L_A + 1)} \right]^4 \right\}^2 (5L_A)^{1/3}, \qquad (7)$$

established to predict a variety of luminance-dependent appearance effects in CIECAM02 and earlier models. Since the function has been established and understood, it was also adopted for the early stages of iCAM. However, the manner in which the F_L factor is used in CIECAM02 and iCAM are quite different.

The next stage of the model is to convert from RGB signals (roughly analogous to cone signals in the human visual system) to opponent-color signals (light-dark, red-green, and yellow-blue; analogous to higher level encoding in the human visual system) that are necessary for constructing a uniform perceptual color space and correlates of various appearance attributes. In choosing this transformation, simplicity, accuracy, and applicability to image processing were the main considerations. The color space chosen was the IPT space previously published by Ebner and Fairchild.[35] The IPT space was derived specifically for image processing applications to have a relatively simple formulation, and specifically, to have a hue-angle component with good prediction of constant perceived hue (important in gamut mapping applications). More recent work on perceived hue has validated the applicability of the IPT space. The transformation from RGB to the IPT opponent space is far simpler than the transformations used in CIECAM02. The process, expressed in Eqs. (9)–(12), involves a linear transformation to a different cone response space (a different RGB), application of power function nonlinearities, and then a final linear transformation to the IPT opponent space (I is light-dark, P is red-green, and T is yellow-blue).

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.4002 & 0.7075 & -0.0807 \\ -0.2280 & 1.1500 & 0.0612 \\ 0.0 & 0.0 & 0.9184 \end{bmatrix} \begin{bmatrix} X_{D65} \\ Y_{D65} \\ Z_{D65} \end{bmatrix}, \qquad (8)$$

$$L' = L^{0.43}; \quad L \geq 0$$
$$L' = -|L|^{0.43}; \quad L \leq 0, \qquad (9)$$

$$M' = M^{0.43}; \quad M \geq 0$$
$$M' = -|M|^{0.43}; \quad M \leq 0, \qquad (10)$$

$$S' = S^{0.43}; \quad S \geq 0$$
$$S' = -|S|^{0.43}; \quad S \leq 0, \qquad (11)$$

$$\begin{bmatrix} I \\ P \\ T \end{bmatrix} = \begin{bmatrix} 0.4000 & 0.4000 & 0.2000 \\ 4.4550 & -4.8510 & 0.3960 \\ 0.8056 & 0.3572 & -1.1628 \end{bmatrix} \begin{bmatrix} L' \\ M' \\ S' \end{bmatrix}. \qquad (12)$$

The power function nonlinearities in the IPT transformation are a critical aspect of the iCAM model. First, they are necessary to predict response compression that is prevalent in most human sensory systems. This response compression helps to convert from signals that are linear in physical metrics (e.g., luminance) to signals that are linear in perceptual dimensions (e.g., lightness). The CIECAM02 model uses a hyperbolic nonlinearity for this purpose. The behavior is that of a power function over the practical ranges of luminance levels encountered. Second, and a key component of iCAM, the exponents are modulated according to the luminance of the image (low-pass filtered) and the surround. This is essentially accomplished by multiplying the base exponent in the IPT formulation by the image-wise computed F_L factors with appropriate normalization. These modulations of the IPT exponents allow the iCAM model to be used for predictions of the Hunt, Stevens, and Bartleson/Breneman effects mentioned before. They also happen to enable the tone mapping of high-dynamic-range images into low-dynamic range display systems in a visually meaningful way.

For image difference and image quality predictions, it is also necessary to apply spatial filtering to the image data to eliminate any image variations at spatial frequencies too high to be perceived. For example, the dots in a printed halftone image are not visible if the viewing distance is sufficiently large. This computation is dependent on viewing distance and is based on filters derived from human contrast sensitivity functions. Since the human contrast sensitivity functions vary for luminance (bandpass with sensitivity to high frequencies) and chromatic (low pass) information, it is appropriate to apply these filters in an opponent space. Thus in image quality applications of iCAM, spatial filters are applied in the IPT space. Since it is appropriate to apply spatial filters in a linear signal space, they are applied in a linear version of IPT prior to conversion into the nonlinear version of IPT for appearance predictions. Johnson and Fairchild have recently discussed some of the important considerations for this type of filtering in image difference applications, and have specified the filters used based on available visual data.[18,43] Since the spatial filtering effectively blurs the image data, it is not desirable for image rendering applications in which observers might view the images more closely than the specified viewing distance. The result would be a blurrier image than the original. The contrast sensitivity functions used to define spatial filters currently used for image difference computations are given in Eq. (13) for the luminance I channel and Eq. (14) for the chromatic P and T channels,[10] respectively.

$$\mathrm{csf}_{\mathrm{lum}}(f) = a \cdot f^c \cdot \exp(-b \cdot f), \qquad (13)$$

$$\mathrm{csf}_{\mathrm{chrom}}(f) = a_1 \cdot \exp(-b_1 \cdot f^{c_1}) + a_2 \cdot \exp(-b_2 \cdot f^{c_2}). \qquad (14)$$

Equations (13) and (14) were derived from fits to a collection of experimental data.[10] The parameters a, b, and c in Eq. (13) are set to 75, 0.2, and 0.8, respectively, for the luminance CSF, applied to the I channel. In Eqs. (13) and (14), spatial frequency f is defined in terms of cycles per degree of visual angle (CPD). For the red-green chromatic CSF, applied to the P dimension, the parameters (a_1, b_1, c_1, a_2, b_2, c_2) in Eq. (14) are set to (109.14, -0.00038, 3.424, 93.60, -0.00367, 2.168). For the blue-yellow chromatic CSF, applied to the T dimension, they are set to (7.033, 0.000004, 4.258, 40.69, -0.10391, 1.6487).

It is only appropriate to apply these spatial filters when the goal is to compute perceived image differences (and ultimately image quality). This is an important distinction between spatially localized adaptation (good for rendering and image quality metrics) and spatial filtering (good for image quality metrics, bad for rendering). In image quality applications, the spatial filtering is typically broken down into multiple channels for various spatial frequencies and orientations. For example, Daly,[17] Lubin,[22] and Pattanaik *et al.*[40] describe such models. More recent results suggest that while such multiscale and multiorientation filtering might be critical for some threshold metrics, it is often not necessary for data derived from complex images and for suprathreshold predictions of perceived image differences (one of the main goals of iCAM).[10,44,45] Thus, to preserve the simplicity and ease of use of the iCAM model, single-scale spatial filtering with anisotropic filters was adopted.

Once the IPT coordinates are computed for the image data, a simple coordinate transformation from rectangular to cylindrical coordinates is applied to obtain image-wise predictors of lightness (J), chroma (C), and hue angle (h), as shown in Eqs. (15), (16), and (17). Differences in these dimensions can be used to compute image difference statistics and those used to derive image quality metrics. The overall Euclidean difference in IPT is referred to as Δ Im (Eq. 20) for image difference, to distinguish it from a traditional color difference metric ΔE, which includes no spatial filtering. In some instances, correlates of the absolute appearance attributes of brightness (Q) and colorfulness (M) are required. These are obtained by scaling the relative attributes of lightness and chroma with the appropriate function of F_L (based on CIECAM02), derived from the image-wise luminance map, as shown in Eqs. (18) and (19).

$$J = I, \tag{15}$$

$$C = \sqrt{P^2 + T^2}, \tag{16}$$

$$h = \tan^{-1}\left(\frac{P}{T}\right), \tag{17}$$

$$Q = \sqrt[4]{F_L}J, \tag{18}$$

$$M = \sqrt[4]{F_L}C, \tag{19}$$

$$\Delta \text{Im} = (\Delta I^2 + \Delta P^2 + \Delta T^2)^{1/2}. \tag{20}$$

For image rendering applications, it is necessary to take the computed appearance correlates (JCh) and then render them to the viewing conditions of a given display. The display viewing conditions set the parameters for the inversion of the IPT model and the chromatic adaptation transform (all for an assumed spatially uniform display adaptation typical of low-dynamic-range output media). This inversion allows the appearance of original scenes or images from disparate viewing conditions to be rendered for the observer viewing a given display. One important application of such rendering is the display of high-dynamic-range (HDR) image data on typical displays.

3 Modular Image Difference Model

A framework for a color image difference metric has recently been described.[10,43] In this work, the modular image difference metric is incorporated into the iCAM appearance model to address both image appearance and differences/quality within a single model. The image difference framework was designed to be modular in nature, to allow for flexibility and adaptation. The framework itself is based on the S-CIELAB spatial extension to the CIELAB color space. S-CIELAB merges traditional color difference equations with spatial properties of the human visual system. This was accomplished as a spatial filtering preprocessing, before a pixel-by-pixel color difference calculation.[9]

The modular framework further extends this idea by adding several processing steps, in addition to the spatial filtering. These processing steps are contained in independent modules, so they can be tested and refined. Several modules have been defined,[45] and include spatial filtering, adaptation, and localization, as well as local and global contrast detection. Figure 2 shows a general flowchart with several distinct modules. These modules and their origins are described briefly in the following.

3.1 Spatial Filtering

The behavior of the human visual system in regards to spatially complex stimuli has been well studied over the years, dating back to the seminal work of Campbell and Robson[46] and Mullen.[47] Summaries of current knowledge and techniques for quantifying spatial vision can be found in several books.[48–50] The contrast sensitivity function describes this behavior in relation to spatial frequency. Essentially, the CSF is described in a postretinal opponent color space, with a bandpass nature for the luminance channel and low-pass nature for the chrominance channels. S-CIELAB uses separable convolution kernels to approximate the CSF, and modulate image details that are imperceptible. More complicated contrast sensitivity functions, that include both modulation and frequency enhancement, were discussed in detail by Johnson and Fairchild.[43] Other models with similar features include the previously mentioned Lubin,[21] Daly,[17] MOM,[40] S-CIELAB,[9] and spatial ATD[34] models. Other relevant discussions and models can be found in the work of Li,[51] Taylor *et al.*,[52,53] and Brill's extension of the Lubin/Sarnoff model.[54]

3.2 Spatial Frequency Adaptation

The contrast sensitivity function in this framework serves to modulate spatial frequencies that are not perceptible, and enhance certain frequencies that are most perceptible. Gen-
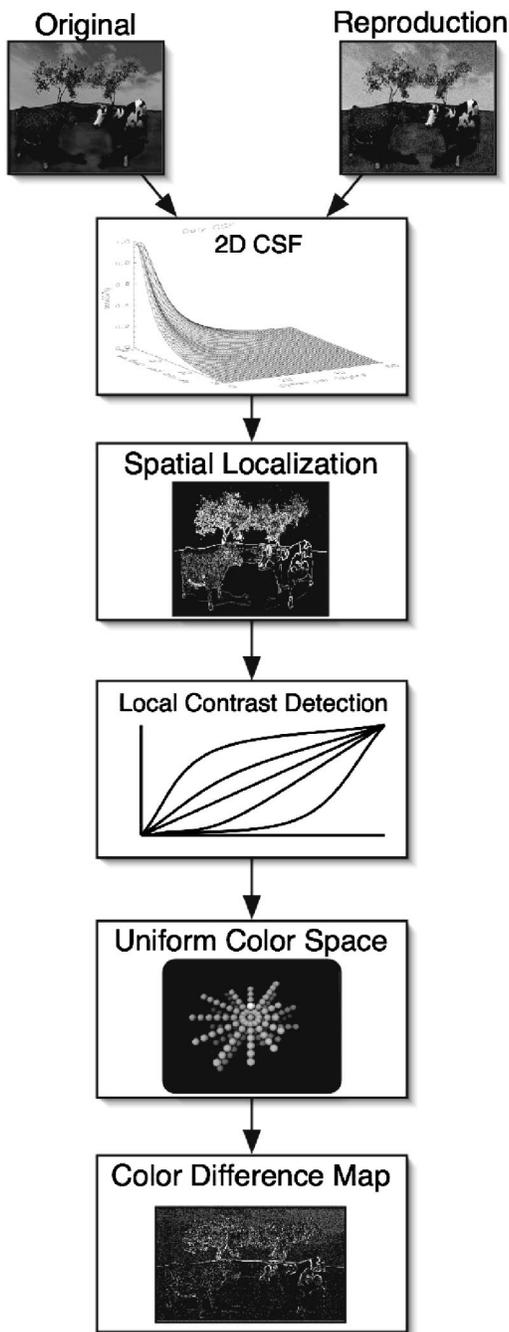
Fig. 2 Flowchart of a modular image difference metric.

erally, CSFs are measured using simple grating stimuli, with care taken to avoid spatial frequency adaptation. Spatial frequency adaptation essentially decreases sensitivity to certain frequencies based on information present in the visual field. An early and classic description of spatial frequency adaptation was published by Blakemore and Campbell.[55] It should be noted that a multiscale, or multichannel, spatial vision model is not required to predict spatial frequency adaptation. Instead, all that is required is that the CSF functions be allowed to change shape as a function of adaptation (clearly indicating multiscale mechanisms in the human visual system not necessary for practical modeling).

Since spatial frequency adaptation cannot be avoided in real world viewing conditions (as it often is in the contrived psychophysical stimuli used to study spatial vision), several models of spatial frequency adaptation have been described for practical applications.[43] These models alter the nature of the CSF based on either assumptions of the viewing conditions, or based on the information contained in the images themselves.

3.3 Spatial Localization

The bandpass and low-pass contrast sensitivity serve to modulate high-frequency information, including high-frequency edges. The human visual system is generally acknowledged to be very adept at detecting edges. To accommodate this behavior, a module of spatial localization has been developed. This module can be as simple as an image processing edge-enhancing kernel, although that kernel must change as a function of viewing distance. Alternatively, the CSF can be modified to boost certain high-frequency information. The formulation and utility of edge-detection algorithms in vision applications has been well described by Marr.[56]

3.4 Local Contrast Detection

This module serves to detect local and global contrast changes between images. The utility of such processing in real visual systems has been described by Tolhurst and Heeger.[57] The current implementation is based on the nonlinear mask-based local contrast enhancement described by Moroney.[58] Essentially, a low-pass image mask is used to generate a series of tone-reproduction curves. These curves are based on the global contrast of the image, as well as the relationship between a single pixel and its local neighborhood.

3.5 Color Difference Map

The output of the modular framework is a map of color differences · Im, corresponding to the perceived magnitude of error at each pixel location. This map can be very useful for determining specific causes of error, or for detecting systematic errors in a color imaging system. Often, it is useful to reduce the error map into a more manageable dataset. This can be accomplished using image statistics, so long as care is taken. Such statistics can be image mean, max, median, or standard deviation. Different statistics might be more valuable than others, depending on the application, as perhaps the mean error better describes overall difference, while the max might better describe threshold differences.

4 Image Appearance Applications (Rendering)

Figure 3 illustrates implementation of the iCAM framework required to complete an image rendering process necessary for HDR image tone mapping. The components essential in this process are the inversion of the IPT model for a single set of spatially constant viewing conditions (the display), and the establishment of spatial filters for the adapting stimuli used for local luminance adaptation and modulation of the IPT exponential nonlinearity. While the derivation of optimal model settings for HDR image ren-
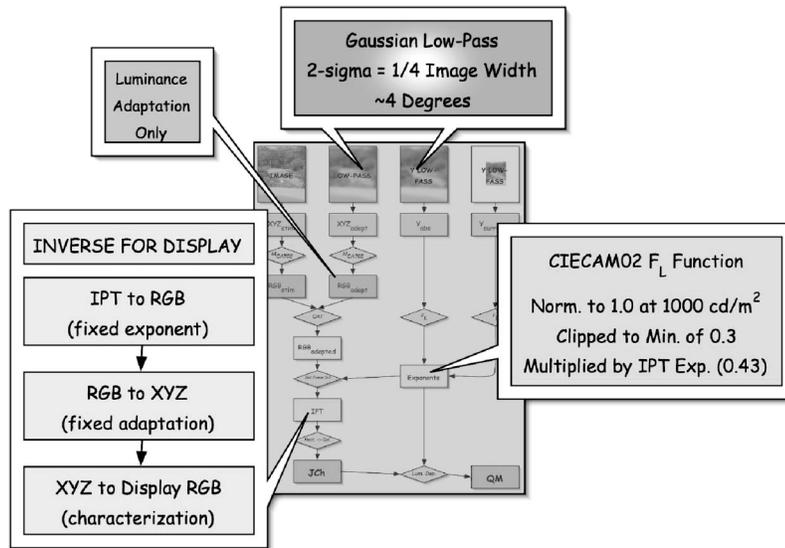
Fig. 3 Implementation of iCAM for tone mapping of HDR images.

dering is still underway, quite satisfactory results have been obtained using the settings outlined in Fig. 3.

The iCAM model has been successfully applied to prediction of a variety of color appearance phenomena, such as chromatic adaptation (corresponding colors), color appearance scales, constant hue perceptions, simultaneous contrast, crispening, spreading, and image rendering.[25]

Since iCAM uses the same chromatic adaptation transform as CIECAM02, it performs identically for situations in which only a change in state of chromatic adaptation is present (i.e., change in white point only). CIE TC8-01 has worked very hard to arrive at this adaptation transform, and it is clear that no other model currently exists with better performance, although there are several with equivalent performance. Thus the chromatic adaptation performance of iCAM is as good as possible at this juncture.

The appearance scales of iCAM are identical to the IPT scales for the reference viewing conditions. The IPT space has the best available performance for constant hue contours, and thus this feature is retained in iCAM. This feature makes accurate implementation of gamut mapping algorithms far easier in iCAM than in other appearance spaces. In addition, the predictions of lightness and chroma

in iCAM are very good and comparable with the best color appearance models in typical viewing conditions. The brightness and colorfulness scales will also perform as well as any other model for typical conditions. In more extreme viewing conditions, the performance of iCAM and other models will begin to deviate. It is in these conditions that the potential strengths of iCAM will become evident. Further visual data must be collected to evaluate the model's relative performance in such situations.

The color difference performance of iCAM will be similar to that of CIELAB, since the space is very similar under the reference viewing conditions. Thus, color difference computations will be similar to those already commonly used, and the space can be easily extended to have a more accurate difference equation following the successful format of the CIE94 equations. (Following the CIEDE2000 equations in iCAM is not recommended, since they are extremely complex and fitted to particular discrepancies of the CIELAB space, such as poor constant hue contours.)

Simultaneous contrast (or induction) causes a stimulus to shift in appearance away from the color of the background in terms of opponent dimensions. Figure 4 illustrates a stimulus that exhibits simultaneous contrast in
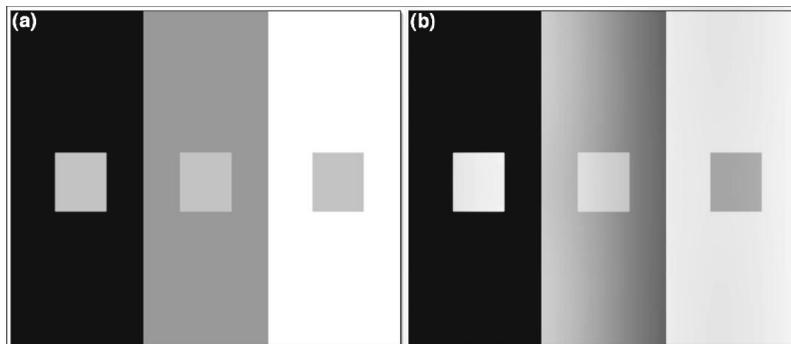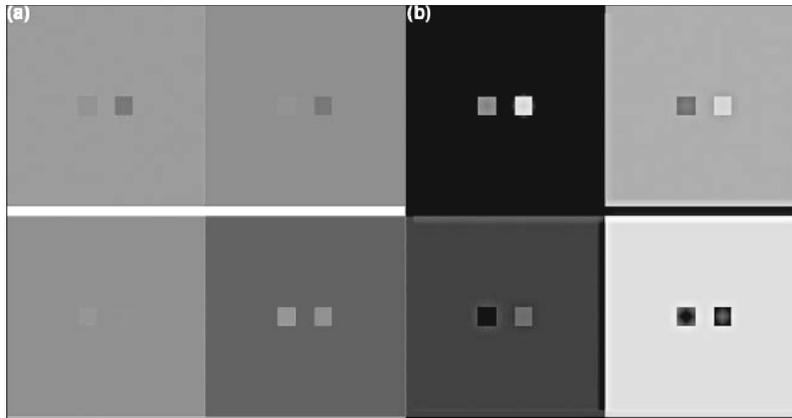

Fig. 4 (a) Original stimulus and (b) iCAM lightness *J* image illustrating the prediction of simultaneous contrast.

Fig. 5 (a) Original stimulus and (b) iCAM chroma *C* image illustrating the prediction of chroma crispening. Original image from www.hpl.hp.com/personal/Nathan_Moroney/.

lightness (the gray square is physically identical on all three backgrounds) and its prediction by iCAM, as represented by the iCAM lightness predictor. This prediction is facilitated by the local adaptation features of iCAM.

Crispening is the phenomenon whereby the color differences between two stimuli are perceptually larger when viewed on a background that is similar to the stimuli. Figure 5 illustrates a stimulus that exhibits chroma crispening, and its prediction by the iCAM chroma predictor. This prediction is also facilitated by the local adaptation features of iCAM.

Spreading is a spatial color appearance phenomenon in which the apparent hue of spatially complex image areas appears to fill various spatially coherent regions. Figure 6 provides an example of spreading in which the red hue of the annular region spreads significantly from the lines to the full annulus. The iCAM prediction of spreading is illustrated through reproduction of the hue prediction. The prediction of spreading in iCAM is facilitated by spatial filtering of the stimulus image.

One of the most interesting and promising applications of iCAM is to the rendering of HDR images to low-dynamic-range display systems. HDR image data are quickly becoming more prevalent. Historically, HDR images were obtained through computer graphics simulations computed with global illumination algorithms (e.g., ray tracing or radiosity algorithms) or through the calibration and registration of images obtained through multiple exposures. Real scenes, especially those with visible light sources, often have luminance ranges of up to six orders of magnitude. More recently, industrial digital imaging systems have become commercially available that can more easily capture HDR image data. It is also apparent that consumer digital cameras will soon be capable of capturing greater dynamic ranges. Unfortunately, display and use of such data are difficult and will remain so, since even the highest quality displays are generally limited in dynamic range to about two orders of magnitude. One approach is to interactively view the image and select areas of interest to be viewed optimally within the display dynamic range. This is only applicable to computer displays and not appropriate for pictorial imaging and printed output. Another limitation is the need for capability to work with greater than 24-bit (and often floating point) image data. It is desirable to render HDR pictorial images onto a display that can be viewed directly (no interactive manipulation) by the observer, and appear similar to what the observer would perceive if the original scene was viewed. For printed images, it is not just desirable, but necessary. Pattanaik *et al.*[40] review several such HDR rendering algorithms, and it is worth noting that several papers were presented on the
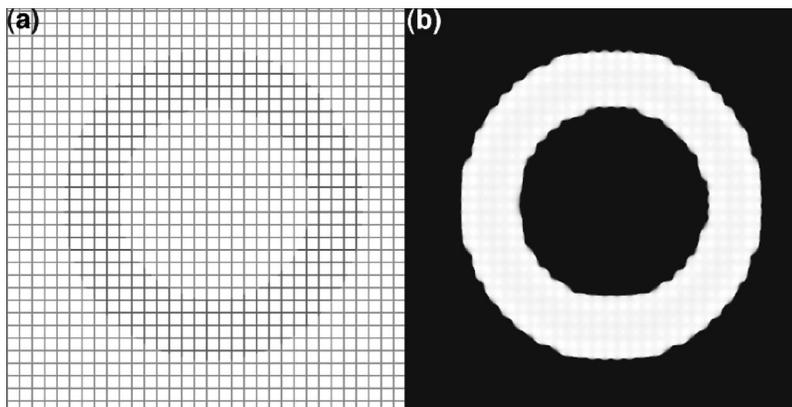
Fig. 6 (a) Original stimulus and (b) iCAM hue *h* image illustrating the prediction of spreading.

Fig. 7 Three HDR images from www.debevec.org. The leftmost column illustrates linear rendering of the image data, the middle column illustrates manually optimized power function transformations, and the rightmost column represents the automated output of the iCAM model implemented for HDR rendering (see Fig. 3).

topic at SIGGRAPH 2002,[59–61] illustrating continued interest in the topic.

Since iCAM includes spatially localized adaptation and spatially localized contrast control, it can be applied to the problem of HDR image rendering. This is not surprising, since the fundamental problem in HDR rendering is to reproduce the appearance of an HDR image or scene on a low-dynamic-range display. Since the encoding in our visual system is of a rather low dynamic range, this is essentially a replication of the image appearance processing that goes on in the human observer and is being modeled by iCAM. Figure 7 illustrates application of the iCAM model to HDR images obtained from Debevec (see www.debevec.org). The images in the left column of Fig. 7 are linear renderings of the original HDR data normalized to the maximum presented, simply to illustrate how the range of the original data exceeds a typical 24-bit (8-bits per RGB channel) image display. For example, the memorial image data (top row) have a dynamic range covering about six orders of magnitude, since the sun was behind one of the stained-glass windows. The middle column of images represents a typical image processing solution to rendering the data. One might consider a logarithmic transformation of the data, but that would do little to change the rendering in the first column. Instead, the middle column was generated interactively by finding the optimum power function transformation (also sometimes referred to as gamma correction; note that the linear images in the first column are already gamma corrected). For these images, transformations with exponents, or gammas, of approximately 1/6 (as opposed to 1/1.8 to 1/2.2 for typical displays) were required to make the image data in the shadow areas visible. While these power function transformations do make more of the image data visible, they required user interaction, tend to wash out the images in a way not consistent with the visual impression of the scenes, and introduce potentially severe quantization artifacts in shadow regions. The rightmost column of images shows the output of the iCAM model with spatially localized adaptation and contrast control (as shown in Fig. 3). These images both render the dynamic range of the scene to make shadow areas visible and retain the colorfulness of the scene. The resulting iCAM images are quite acceptable as reproductions of the HDR scenes (equivalent to the result of dodging and burning historically done in photographic printing). It is also noteworthy that the iCAM-rendered images were all computed with an automated algorithm mimicking human perception with no user interaction.

5 Image Quality Applications (Difference Perceptibility)

A slightly different implementation of iCAM is required for image quality applications to produce image maps representing the magnitude of perceived differences between a pair of images. In these applications, viewing-distance-dependent spatial filtering is applied in a linear IPT space, and then differences are computed in the normal nonlinear IPT space. Euclidean summations of these differences can be used as an overall image difference map, and then various summary statistics can be used to predict different attributes of image difference and quality. This process is outlined in Fig. 8 and detailed in Johnson and Fairchild.[45]

Image quality metrics can be derived from image difference metrics that are based on normal color difference formulas applied to properly spatially filtered images. This approach has been used to successfully predict various types of image quality data.[10] Figure 9 illustrates the prediction of perceived sharpness[62] and contrast[63] differences in images through a single summary statistic (mean image difference). This performance is equivalent to, or better than, that obtained using other color spaces optimized for the task.[10]

The contrast results in Fig. 9(a) were obtained by asking observers to scale perceived image contrast for a collection of images of various content, subjected to a variety of transformations.[63] The resulting interval scale (average data) is plotted as perceived contrast in Fig. 9(a), and the model prediction of image difference from the original (arbitrarily selected) is compared with it. Ideally, the data would follow a V shape with two line segments of equal absolute slope on either side of the origin. The perceived contrast data are well predicted by the iCAM image difference.

The perceived sharpness results in Fig. 9(b) were obtained in a similar manner using a significantly larger number of image manipulations and content.[62] Observers were simply asked to scale perceived sharpness, and the results were converted to an interval scale, again with the original image as an arbitrary zero point. There is greater variability in these data, but it can be seen in Fig. 9(b) that the results are again well predicted by a fairly simple mean image difference metric.

6 Conclusions and Future Directions

Advances in imaging and computing technologies, along with increased knowledge of the function and performance of the human visual system, have allowed for the integration of models of color, spatial, and temporal vision to create a new type of color appearance model, referred to as an
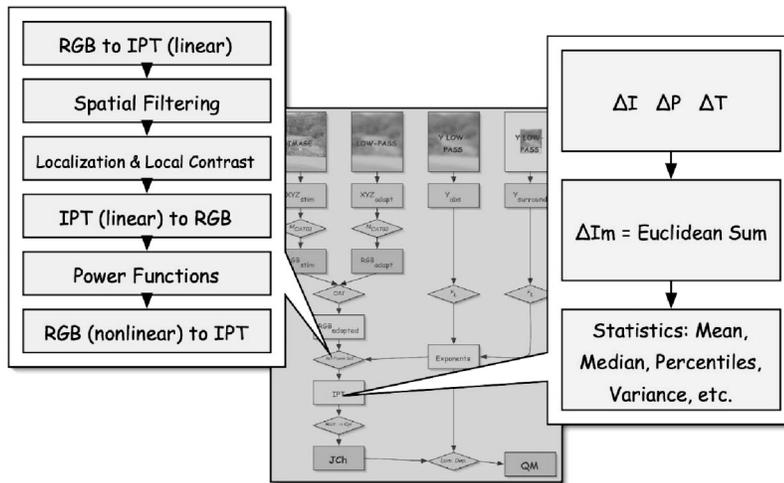
Fig. 8 Implementation of iCAM for image difference and image quality metrics.

image appearance model. Such models show promise in a variety of applications, ranging from image difference and image quality metrics to the rendering of image data. This work describes the framework of one example of an image
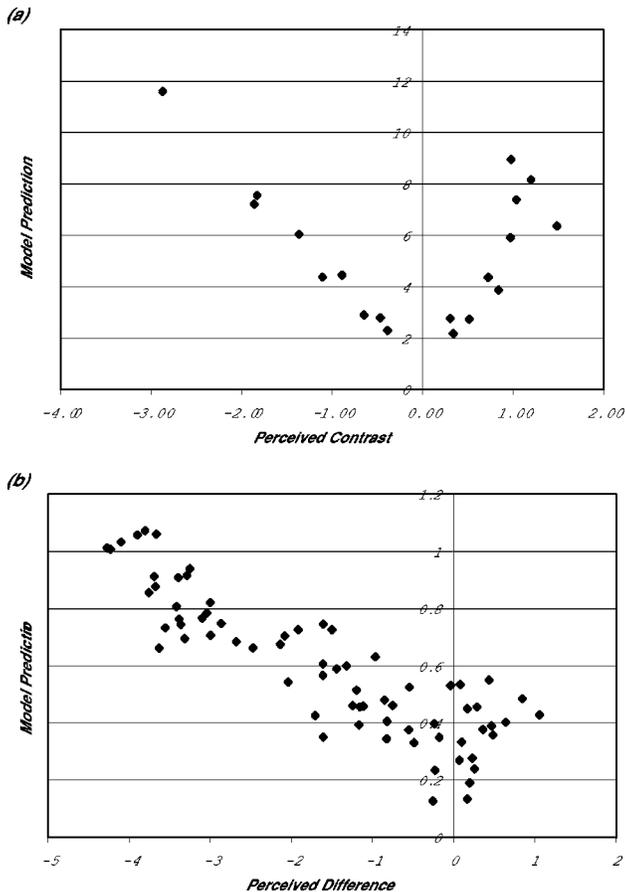

Fig. 9 iCAM image differences as a function of (a) perceived image contrast and (b) perceived image sharpness for a variety of image transformations. (Note that desired predictions are V-shaped data distributions, since the perceptual differences are signed and the calculated differences are unsigned.)

appearance model, referred to as iCAM, and illustrates its applicability to HDR image tone mapping and image quality metrics. Recently, initial efforts have been made to incorporate psychophysical data on the time course of chromatic adaptation[15] to extend the model to video appearance and quality applications.[64] Future efforts will be directed at completion of the spatio-temporal filters required for video difference metrics, the collection of more psychophysical data on image and video appearance and differences, and the formulation of specific iCAM algorithms for various applications. The iCAM model is not proprietary. Source code and updates are freely available at www.cis.rit.edu/mcsl/iCAM for those interested in evaluating the model and potentially suggesting improvements.

References

1. R. W. G. Hunt, *The Reproduction of Colour*, 5th ed., Fountain Press, England (1995).
2. O. H. Schade, "Optical and photoelectric analog of the eye," *J. Opt. Soc. Am.* **46**, 721–739 (1956).
3. W. D. Wright, "50 years of the 1931 CIE standard observer for colorimetry," *AIC Color 81*, Paper A3 (1981).
4. R. S. Berns, "A generic approach to color modeling," *Color Res. Appl.* **22**, 318–325 (1997).
5. M. D. Fairchild, "Some hidden requirements for device-independent color imaging," *SID Intl. Symp.* pp. 865–868 (1994).
6. G. J. Braun and M. D. Fairchild, "General-purpose gamut-mapping algorithms: Evaluation of contrast-preserving rescaling functions for color gamut mapping," *J. Imaging Sci. Technol.* **44**, 343–350 (2000).
7. G. M. Johnson and M. D. Fairchild, "Full-spectral color calculations in realistic image synthesis," *IEEE Comput. Graphics Appl.* **19**(4), 47–53 (1999).
8. M. R. Luo, G. Cui, and B. Rigg, "The development of the CIE 2000 colour difference formula," *Color Res. Appl.* **26**, 340–350 (2001).
9. X. M. Zhang and B. A. Wandell, "A spatial extension to CIELAB for digital color image reproduction," *Proc. SID Symp.* pp. 731–734 (1996).
10. G. M. Johnson and M. D. Fairchild, "Darwinism of color image difference models," *Proc. of IS&T/SID 9th Color Imag. Conf.*, pp. 108–112 (2001).
11. CIE, "The CIE 1997 interim colour appearance model (simple version)," CIECAM97s, *CIE Pub. 131* (1998).
12. N. Moroney, M. D. Fairchild, R. W. G. Hunt, C. J. Li, M. R. Luo, and T. Newman, "The CIECAM02 color appearance model," *IS&T/SID 10th Color Imag. Conf.*, pp. 23–27 (2002).
13. M. D. Fairchild, *Color Appearance Models*, Addison-Wesley, Reading, MA (1998).
14. K. M. Braun and M. D. Fairchild, "Testing five color appearance models for changes in viewing conditions," *Color Res. Appl.* **22**, 165–174 (1997).
15. M. D. Fairchild and L. Reniff, "Time-course of chromatic adaptation

for color-appearance judgements,'' *J. Opt. Soc. Am. A* **12**, 824–833 (1995).

16. O. Rinner and K. R. Gegenfurtner, ''Time course of chromatic adaptation for color appearance discrimination,'' *Vision Res.* **40**, 1813–1826 (2000).

17. S. Daly, ''The visible differences predictor: An algorithm for the assessment of image fidelity,'' in *Digital Images and Human Vision*, A. Watson, Ed., pp. 179–206, MIT Press Cambridge, MA (1993).

18. G. M. Johnson and M. D. Fairchild, ''A top down description of S-CIELAB and CIEDE2000,'' *Color Res. Appl.* (in press).

19. Sarnoff Corporation, ''JND: A human vision system model for objective picture quality measurement,'' Sarnoff Technical Report from www.jndmetrix.com (2001).

20. ATIS, ''Objective perceptual video quality measurement using a JND-based full reference technique,'' Alliance for Telecommunications Industry Solutions Technical Report T1.TR.PP.75-2001 (2001).

21. J. Lubin, ''The use of psychophysical data and models in the analysis of display system performance,'' in *Digital Images and Human Vision*, A. Watson, Ed., pp. 163–178, MIT Press, Cambridge, MA (1993).

22. J. Lubin, ''A visual discrimination model for imaging system design and evaluation,'' in *Vision Models for Target Detection and Recognition*, E. Peli, Ed., pp. 245–283, World Scientific, Singapore (1995).

23. A. B. Watson, ''Toward a perceptual video quality metric,'' *Proc. SPIE* **3299**, 139–147 (1998).

24. A. B. Watson, J. Hu, and J. F. McGowan, ''DVQ: A digital video quality metric based on human vision,'' *J. Electron. Imaging* **10**, 20–29 (2001).

25. M. D. Fairchild and G. M. Johnson, ''Meet iCAM: A next-generation color appearance model,'' *IS&T/SID 10th Color Imag. Conf.*, pp. 33–38 (2002).

26. M. D. Fairchild, ''Image quality measurement and modeling for digital photography,'' *International Congress on Imaging Science '02*, pp. 318–319 (2002).

27. E. H. Land, ''Recent advances in retinex theory,'' *Vision Res.* **26**, 7–21 (1986).

28. E. H. Land and J. J. McCann, ''Lightness and the retinex theory,'' *J. Opt. Soc. Am.* **61**, 1–11 (1971).

29. E. H. Land, ''The retinex,'' *Am. Sci.* **52**, 247–264 (1964).

30. J. J. McCann, S. McKee, and T. Taylor, ''Quantitative studies in retinex theory: A comparison between theoretical predictions and observer responses to 'color Mondrian' experiments,'' *Vision Res.* **16**, 445–454 (1976).

31. B. Funt, F. Ciurea, and J. J. McCann, ''Retinex in Matlab,'' *Proc. of IS&T/SID 8th Color Imag. Conf.,* pp. 112–121 (2000).

32. K. Barnard and B. Funt, ''Analysis and improvement of multi-scale retinex,'' *Proc. 5th IS&T/SID Color Imag. Conf.*, pp. 221–226 (1997).

33. D. H. Brainard and B. A. Wandell, ''Analysis the retinex theory of color vision,'' *J. Opt. Soc. Am. A* **3**, 1651–1661 (1986).

34. E. M. Granger, ''Uniform color space as a function of spatial frequency,'' *Proc. SPIE* **1913**, 449–457 (1993).

35. F. Ebner and M. D. Fairchild, ''Development and testing of a color space (IPT) with improved hue uniformity,'' *IS&T/SID 6th Color Imag. Conf.*, pp. 8–13 (1998).

36. M. D. Fairchild, ''Considering the surround in device-independent color imaging,'' *Color Res. Appl.* **20**, 352–363 (1995).

37. M. D. Fairchild, ''Modeling color appearance, spatial vision, and image quality,'' *Color Image Science: Exploiting Digital Media*, pp. 357–370, Wiley and Sons, New York (2002).

38. J. S. Babcock, J. B. Pelz, and M. D. Fairchild, ''Eye tracking observers during color image evaluation tasks,'' *Proc. SPIE* **5007**, 218–230 (2003).

39. M. A. Webster and J. D. Mollon, ''Adaptation and the color statistics

40. of natural images,'' *Vision Res.* **37**, 3283–3298 (1997).

40. S. N. Pattanaik, J. A. Ferwerda, M. D. Fairchild, and D. P. Greenberg, ''A multiscale model of adaptation and spatial vision for image display,'' *Proc. SIGGRAPH 98*, pp. 287–298 (1998).

41. B. W. Keelan, *Handbook of Image Quality: Characterization and Prediction*, Marcel Dekker, New York (2002).

42. P. G. Engledrum, ''Extending image quality models,'' *Proc IS&T PICS Conf.*, 65–69 (2002).

43. G. M. Johnson and M. D. Fairchild, ''On contrast sensitivity in an image difference model,'' *Proc. IS&T PICS Conf.*, 18–23 (2001).

44. A. B. Watson and C. V. Ramirez, ''A standard observer for spatial vision,'' *Investigative Ophthalmology Vis. Sci.* **41**, S713 (2000).

45. G. M. Johnson and M. D. Fairchild, ''Measuring images: Differences, quality, and appearance,'' *Proc. SPIE* **5007**, 51–60 (2003).

46. F. W. Campbell and J. G. Robson, ''Application of Fourier analysis to the visibility of gratings,'' *Am. J. Physiol.* **197**, 551–566 (1968).

47. K. T. Mullen, ''The contrast sensitivity of human color vision to red-green and blue-yellow chromatic gratings,'' *Am. J. Physiol.* **359**, 381–400 (1985).

48. R. L. DeValois and K. K. DeValois, *Spatial Vision*, Oxford University Press, Oxford (1988).

49. *Visual Science and Engineering: Models and Applications*, D. H. Kelly, Ed., Marcel Dekker, New York (1994).

50. B. A. Wandell, *Foundations of Vision*, Sinauer Associates Inc., Sunderland, MA (1995).

51. B. Li, G. W. Meyer, and R. V. Klassen, ''A comparison of two image quality models,'' *Proc. SPIE* **3299**, 98–109 (1998).

52. C. C. Taylor, Z. Pizlo, J. P. Allebach, and C. A. Bouman, ''Image quality assessment with a Gabor pyramid model of the human visual system,'' *Proc. SPIE* **3016**, 58–69 (1997).

53. C. C. Taylor, Z. Pizlo, and J. P. Allebach, ''Perceptually relevant image fidelity,'' *Proc. SPIE* **3299**, 110–118 (1998).

54. M. H. Brill, ''Color management: New roles for color transforms,'' *IS&T/SID 5th Color Imaging Conf.*, pp. 78–82 (1997).

55. C. Blakemore and F. W. Campbell, ''On the existence of neurons in the human visual system selectively sensitive to the orientation and size of retinal images,'' *Am. J. Physiol.* **203**, 237–260 (1969).

56. D. Marr, *Vision*, Freeman, New York (1982).

57. D. J. Tolhurst and D. J. Heeger, ''Comparison of contrast-normalization and threshold models of the responses of simple cells in cat striate cortex,'' *Visual Neurosci.* **14**, 293–309 (1997).

58. N. Moroney, ''Local color correction using non-linear masking,'' *Proc. IS&T/SID 8th Color Imaging Conf.*, pp. 108–111 (2000).

59. R. Fattal, D. Lischinski, and M. Werman, ''Gradient domain high dynamic range compression,'' *Proc. SIGGRAPH 2002*, pp. 249–256 (2002).

60. F. Durand and J. Dorsey, ''Fast bilateral filtering for the display of high-dynamic-range images,'' *Proc. SIGGRAPH 2002*, pp. 257–266 (2002).

61. E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, ''Gradient domain high dynamic range compression,'' *Proc. SIGGRAPH 2002*, pp. 267–276 (2002).

62. G. M. Johnson and M. D. Fairchild, ''Sharpness rules,'' *Proc. IS&T/SID 8th Color Imaging Conf.*, pp. 24–30 (2000).

63. A. J. Calabria and M. D. Fairchild, ''Compare and contrast: Perceived contrast of color images,'' *Proc. IS&T/SID 10th Color Imaging Conf.*, 17–22 (2002).

64. M. D. Fairchild and G. M. Johnson, ''Image appearance modeling,'' *Proc. SPIE* **5007**, 149–160 (2003).

Biographies and photographs of the authors not available.