Wide-gamut image capture

Charles Poynton
Simon Fraser University, CANADA
cpoynton@sfu.ca

Abstract: *Colour gamut* refers to the range of colours that can be reproduced by an imaging system. The definition of gamut is quite clear for displays and for hard-copy printing. Colour image science experts disagree, however, on the definition – or even applicability of the concept – of gamut for cameras. I disagree that there is any meaningful concept of "capture gamut." In this note, I review trichromacy and metamerism and discuss various gamuts. I conclude that although metamerism is a phenomenon of what I call "31-space," gamut lives in 3-space. With suitable colour signal processing, a 3-channel camera is capable of acquiring wide gamut images.

Introduction

The principles of colour science, and its application in video, are discussed in chapters 21 and 22 of *Digital Video and HDTV Algorithms and Interfaces* ("DVAI") [1].

Typical electronic displays – such as CRT, LCD, PDP, or DMD displays – have additive *RGB* primaries. Owing to the three types of cone photoreceptors in normal human vision, three well-chosen primaries are necessary and sufficient to achieve metameric colour matching for a wide range of colours.

Multispectral refers[1] to a device having a few, or perhaps several, spectral components beyond the three that are necessary for trichromatic capture. *Hyperspectral* refers to a device having more than a handful of spectral components. There is no accepted definition of how many components constitute multispectral or hyperspectral; in my view, a multispectral system has between 4 and 8 spectral components, and a hyperspectral system has 9 or more. Multispectral displays have been demonstrated, but none are commercially deployed. Experimental multispectral and even hyperspectral cameras are in use, but as I write, none are used in commercial pictorial imaging. Apart from highly specialized applications such as remote sensing, or the preservation or reproduction of fine art, I argue that multispectral techniques are not necessary to capture wide-gamut colour.

Subtractive colour (CMY and CMYK) is used in colour photography and motion pictures, and in commercial and consumer printing. Subtractive reproduction is more complicated than additive, owing to the nonlinearity of colour mixture. Although it is theoretically possible to form colour in an electronic display using the subtractive mechanism, no such display has been commercialized. In the remainder of this note I will address just 3-component additive displays.

Individual colour-normal observers have different spectral sensitivities. For purposes of colour engineering the CIE has adopted a statistically-derived *standard observer* that is the basis for measurement and characterization of colour. The standard observer is defined in terms of three weighting functions termed *colour matching functions* (CMFs). Measuring colour involves forming three weighted integrals of the spectral power distribution (SPD) of the light – one for each CMF curve. The three components that result are termed *tristimulus values,*[2] or simply *tristimuli.* The CIE standard tristimuli are *XYZ* components, associated with $\bar{x}$, $\bar{y}$, and $\bar{z}$ spectral responsivities; other components (such as various flavours of *RGB*) can be obtained from *XYZ* through a 3×3 matrix multiplication.

1 Nearly all of the multispectral cameras described in the research literature, and all of the hyperspectral cameras, involve changing filters in time sequence: Such cameras are unsuitable for capturing moving subjects.

2 Absolute luminance has SI units of $cd \cdot m^{-2}$ [nit, nt]. Relative luminance – that is, luminance relative to a reference white luminance – can be considered to be a distinguished tristimulus value that is meaningful on its own. Apart from relative luminance, tristimuli come in sets of three, as the word suggests, and have no units.

It is sufficient for most colour image system engineering purposes to express the colour matching functions of the standard observer at 31 wavelengths, from 400 nm to 700 nm in steps of 10 nm, ordinarily represented in a 31-by-3 (tall) matrix. Tristimuli are formed by taking the matrix product of a 31-element column vector (containing an SPD) placed to the right of this CMF matrix. The matrix product *projects* – or in common language, collapses – the 31 dimensions of spectral space into the 3 dimensions of colour.

Metamerism

The mapping of spectra to tristimuli is many-to-one. All SPDs that produce the same tristimuli are termed *metamers.* The matching of colour of any pair of these spectra is termed a *metameric match* (as opposed to a spectral match). In some applications it is useful to associate a set of three tristimulus values with a preferred or distinguished SPD called the *fundamental metamer;* other SPDs are then ordinary metamers. Any SPD can be mapped into its fundamental metamer through matrix multiplication with Cohen [2] and Kappauf's *matrix **R**,* described in their 1985 paper. Matrix ***R*** has rank 3; for 31-component spectral sampling, it is 31×31. Matrix ***R*** incorporates an illuminant.

Metamerism is both good and bad news. The good news is that three components suffice to reproduce colour of light on its way to the eye. However, the colours of reflective objects or media involve ambient light. When we see an object, the spectral power distribution of the illuminant interacts wavelength-by-wavelength with the spectral reflectance of the object. The extent to which the spectral character of the ambient light is uncontrolled leads to the bad news of metamerism: Colours can and do change depending upon the spectral composition of illumination.

Emissive displays generate light without depending upon ambient illumination, so they do not suffer from metamerism. However, metamerism affects reflective displays, and it affects image capture.

To represent the colour of an object with anything less than a spectral representation – say in 31 components – the dependence upon illumination is implicated. There are many different illuminants. We can use colourimetry to characterize the *colour* of an illuminant, but any representation in just 3 components cannot adequately capture spectral information: No 3-component representation can accurately capture the interaction between the illuminant and an arbitrary object.

In photographic printing, illuminant SPDs and the spectral reflectances or spectral transmittances of photographic material are well-controlled. Providing that the photographic media has three colourants (as is nearly always the case), three components suffice to represent captured colour. Colour reproduction could be characterized in terms of tristimuli related to the spectral sensitivities of human vision. However, for process control reasons it is usual to characterize photographic reproduction using *optical density* quantities that are directly related to the physics of reproduction. Description of colour in this manner is called *densitometric* instead of *colourimetric.*

Having established the context for a discussion of gamut, I will briefly outline cameras, then proceed to the issue of camera metamerism.

Cameras

Colour cameras filter incoming light into spectral bands, then direct filtered light onto sensors.[3] Typically the sensors are identical for all channels; colour response is dominated by the filter characteristics. In nearly all commercial cameras, three bands are separated. (Experimental cameras having up to six bands have been demonstrated.) Two classes of camera are distinguished according to how they accomplish filtering: beamsplitter cameras and mosaic cameras.

- A *beamsplitter camera* uses dichroic filters in the optical path, interposed between the lens and a set of sensors, to separate a single beam of light into three constituent wavelength bands. An image for each wavelength band is incident upon each sensor. Dichroic filters are not absorptive: No light is lost in colour separation.
- A *mosaic camera* uses a single sensor. A few different colour filter materials are deposited

3 Foveon's X3 sensor is an exception: Colours of light are separated by their absorbtion depth in a three-layer photosite.

onto neighboring sensor elements in a spatially periodic pattern. The scheme invented in 1976 by Kodak researcher Bruce Bayer remains the most common pattern today: The *Bayer* pattern tiles R-G-G-B filters in a 2×2 pattern.[4] Mosaic sensors confound spatial detail and colour; "demosaicking" algorithms are necessary.

Sensor spectral sensitivity and spectral transmittance of the lens and other optical components affect overall spectral sensitivity of a colour camera, but the colour separation mechanism dominates.

Camera metamerism

Human vision has colour-matching functions (CMFs); an electronic colour camera is said to have *spectral responsivity functions* (SRFs).

A camera having SRFs identical to the CIE CMFs (or linear combinations of them) is said to meet the *Maxwell-Ives criterion.*[5] A camera having the same $\bar{x}$, $\bar{y}$, and $\bar{z}$ spectral responsivities as the standard observer would deliver *XYZ* components, and could be called an *XYZ* camera – or could be called an imaging colourimeter.

Life with *a* perfectly colourimetric *XYZ* camera would be simple. If fact, experimental cameras have been described and demonstrated – see the paper by Ejaz and his colleagues [3].

However, there are good engineering reasons – such as optimizing signal-to-noise ratio, or allowing reasonably inexpensive optical filters – to use sensitivities different from the CIE CMFs. The signal-to-noise issue derives from the large degree of spectral overlap between the L and M photoreceptors of vision. Reconstruction of additive *RGB* primary components from highly-overlapped sensor SRFs requires large coefficients in the required 3×3 linear matrix, as I explain in *DVAI* [1] (in the section *Noise due to matrixing,* on page 253). The large matrix coefficients incur a significant noise penalty.

4 Sony commercialized a consumer digital still camera (DSC-F828) having a mosaic sensor with four channels: the usual red, green, and blue, and a fourth "emerald" colour (RGB+E). The fourth channel was claimed to improve colour performance, but I am unaware of any published technical information that supports the claim.

5 What I call the *Maxwell-Ives criterion* is sometimes called *Luther-Ives,* or just *Luther.* In my view, Maxwell and Ives mainly deserve the credit.

To the extent that camera spectral sensitivities depart from the CIE CMFs, the camera will "see" colours differently than human vision: a pair of SPDs that we see as two different colours might produce identical sets of camera responses; conversely, a pair of SPDs that we see as identical might produce distinct sets of camera responses. Departure of the camera response from vision's response, as estimated by CIE colourimetry, is known as *camera metamerism.* Camera metamerism is inherent in any system that departs from the CIE CMFs – and practical systems do depart, so metamerism is certain to occur. Where in colour space the metamerism occurs, and its effect, is not obvious; these are matters to be investigated.

Scanner metamerism relates to a similar phenomenon in scanners. Because objects being scanned are usually colour reproductions that have only three colourants, metamerism is easier to avoid or correct than it is for arbitrary scenes. I won't discuss scanner metamerism any further.

In my view, the practical issue of camera metamerism is not yet well understood. Camera gamut also deserves discussion, but in my view it is not productive to confound camera metamerism and camera gamut. For me, the distinction between these topics is that metamerism takes place in spectral domain – call it "31-space" – and gamut is a phenomenon of 3-space.

Optimal colours

The *optimal colours,* first investigated by David MacAdam in 1935 [4], comprise a set of artificial spectral reflectances that produce as wide a gamut as is possible from a diffusely reflecting surface. The optimal colours form a surface bounding the *object-colour solid* (OCS) that is defined as the set of all possible ideal diffuse spectral reflectances. The optimal colours are defined without reference to any illuminant; they are more accurately called optimal *reflectances*. (When lit, they become optimal colours.) Although optimal reflectances are defined without reference to any illuminant, it is common to discuss them in the context of the equi-energy illuminant (CIE Illuminant E).

For n wavelength samples, there are $1/2 \cdot n \cdot (n+1)$ type 1 reflectances and $1/2 \cdot n \cdot (n+1)$ type 2 reflectances.

Figure 1 **Optimal colors** have either Type 1 spectral reflectances (at the top) or Type 2 spectral reflectance (at the bottom). The optimal colors shown here have the same reflectance factor of 20%. These figures are adapted from Francisco Martínez-Verdú and his colleagues.

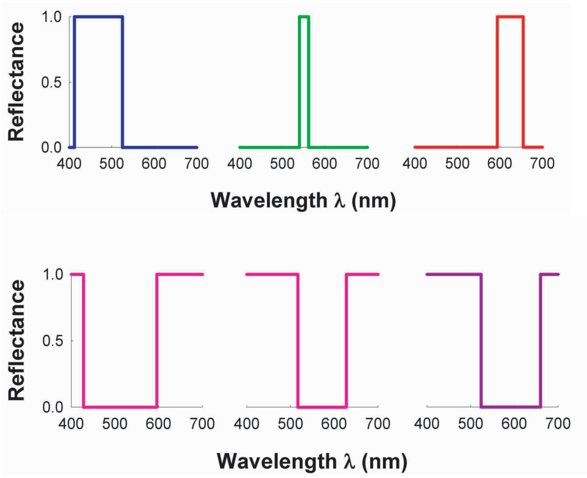

MacAdam proved that optimal colours have just two types of spectral reflectance, both limited to zero reflectance or unit reflectance at each wavelength and having at most two transitions between those values across the visible spectrum. Type 1 spectra are "mountain" shaped, having zero reflectance except for a single spectral peak between λ_1 and λ_2. Type 2 spectra are "valley" shaped, having unity reflectance except for a single notch of zero reflectance between λ_1 and λ_2.

Figure 1 shows six optimal spectral reflectance curves, all at luminance factor of 20%, taken from the paper of Francisco Martínez-Verdú and his colleagues [5]. Verdú uses 41 components 10 nm intervals from 380 nm to 780 nm.[6]

Optimal reflectances are never encountered in practice: Real object surface reflectances never exhibit transitions from zero to unity in an infinitesimal wavelength interval, and never have perfect absorbance. Nonetheless, the optimal reflectances provide a useful analytical tool to explore gamut limits. Significantly, the optimal reflectances have no metamers, so they offer a good way to explore capture gamut without introducing the complications of metamerism.

Numerosity

To estimate the impact of metamerism on the operation of real cameras capturing real scenes, it's important to know something about the frequency of metamerism in natural and synthetic scenes. How many colours, and how many metamers, are encountered?

I have mentioned that MacAdam's optimal colours are unrealistic, because infinitesimally narrow transitions between full reflectance and full absorbance don't occur in nature.

For 31-component spectral sampling, there are 2^{31} – or about two billion – spectral combinations of distinct monochromatic sources, but only 32 times 31 – or 992 – optimal spectral samples. Of the two billion samples, only about a thousand lie on the gamut boundary; the remainder lie within the boundary, and nearly all lie well within.

In 1962, Stiles and Wyszecki published a paper [6] describing a study that they performed to analyze metamers using Monte Carlo techniques, producing a 3-D histogram. Figure 3 reproduces a histogram from that paper. Stiles and Wyszecki conclude that metamers are far more likely to be located within the gamut boundary than near the boundary. (As I mentioned earlier, colours on the gamut boundary – the optimal colours – cannot be metameric.)

Stiles, Wyszecki, and Ohta investigated metamers with spectral constraints making them less "spiky." Their 1977 paper [7] explains use of Fourier techniques to explore non-spiky metamers. These papers, and papers by several researchers following them, confirm that high degrees of metamerism produce tristimuli well within the gamut boundary.

"Spikiness" can arise not only from spiky reflectance but also from spiky illumination:

6 For reasons that are not clear, Verdú uses 1734 reflectances instead of the 1722 that I would expect.

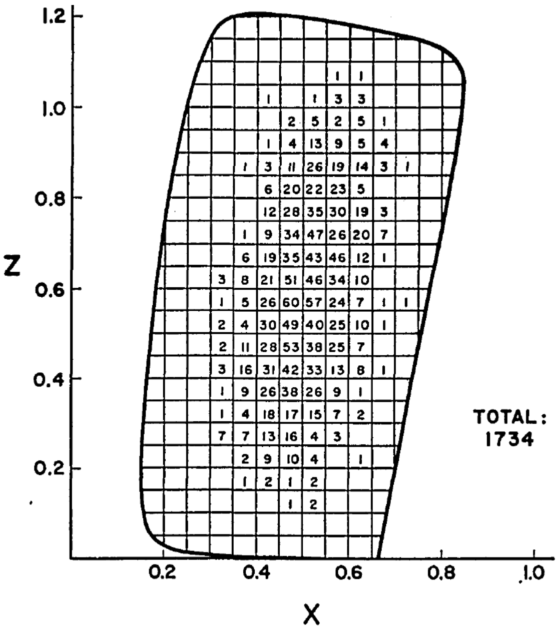

Figure 2 **Stiles' and Wyszecki's histogram** counts metamers produced by the Monte Carlo technique. This figure shows a small cube of *XYZ* space in *a* 2-D slice at luminance (*Y*) values between 0.50 and 0.55. The number in each cell shows the count of metamers lying within the corresponding tristimulus value boundaries. It is evident that most metamers lie well within the gamut boundary.

Mercury-vapor and sodium vapor lamps, for example, have rapid transitions in their SPD curves. However, such lamps are useless in imaging. Many fluorescent lamps are somewhat spiky, but a professional would never except under very severe circumstances capture an image lit by fluorescent lamps. CRT red phosphors have notoriously spiky SPDs, owing to the bicomponent rare-earth phosphor composition. Although the spiky characteristic leads to some difficulties in measurement, CRTs are not used as sources of illumination, so we can discount them as light sources. I conclude that spiky sources do not present serious problems in practice.

Pointer's colours

Mike Pointer, working at Kodak Research in the U.K., collected about two thousand colourimetric samples of real surface reflectances. He published a paper [8] summarizing the CIE $L^*u^*v^*$ and CIE $L^*a^*b^*$ coordinates of colours at the boundary of his set. Pointer plotted his data in a set of 2-D graphs and plots; Figure 3 is Pointer's own representation of gamut as "lightness contours" in CIE [u', v'] chromaticity coordinates. Figure 4 shows my 3-D representation of Pointer's gamut in CIE [L^*, u^*, v^*] coordinates. The gamut of real surface colours is best described as a blob. Many of Pointer's colours are outside of the capability of BT.709/sRGB. It is a goal of wide colour gamut systems to capture and reproduce many of these colours.

Camera capture analysis gamut

Some researchers argue that gamut is limited when a change in the optical stimulus produces no change in sensor output. I disagree with this view. We already have perfectly good words *saturation* (referring to the sensor itself) and *clipping* (referring to signal processing) that express absence of signal change to a changing stimulus. I argue that it is a mistake to confound *gamut* with saturation and clipping.

Peter Centen, lead designer of the Thomson Viper FilmStream camera, has argued that gamut limitation in a camera is not a function of the sensor spectral characteristics, but of signal processing alone – and more specifically, a function of clipping [9]. I agree with his view.

Munsell Color Science Laboratory has, on its web site [10], a column "Ask a color scientist!" One of the answers states, quite unequivocally, "there is no such thing as a camera, or scanner, gamut."

Video cameras, digital still cameras, and digital cinema cameras incorporate signal processing elements to adapt the spectral sensitivities of the sensor to the *RGB* primaries of the assumed display device. The usual signal processing element is a 3×3 "linear" matrix, so named because its action takes place in the

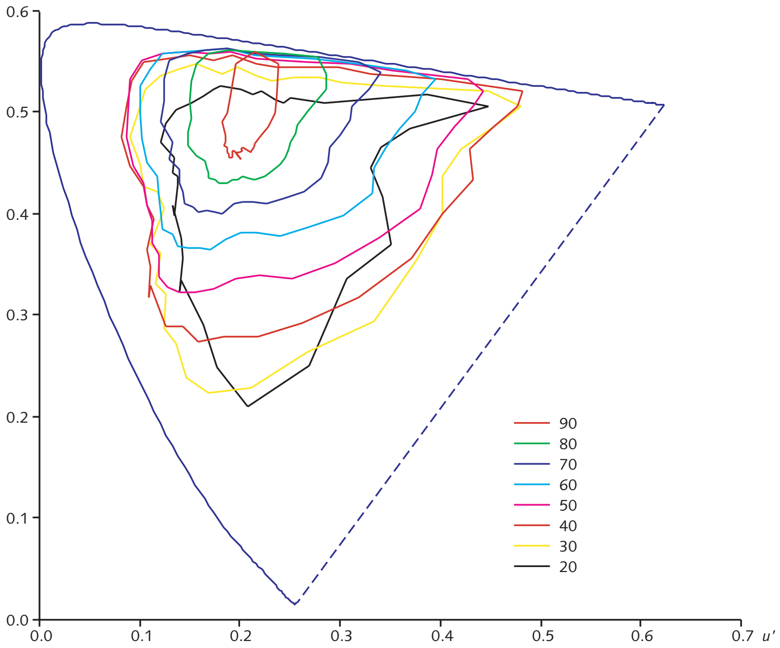

Figure 3 **Pointer's gamut in 2-D** CIE *u'v'* coordinates is plotted as lightness (*L**) contours at the indicated levels.

Figure 4 **Pointer's gamut in 3-D,** here in CIE *L*u*v** coordinates, forms what is best described as *a* blob.

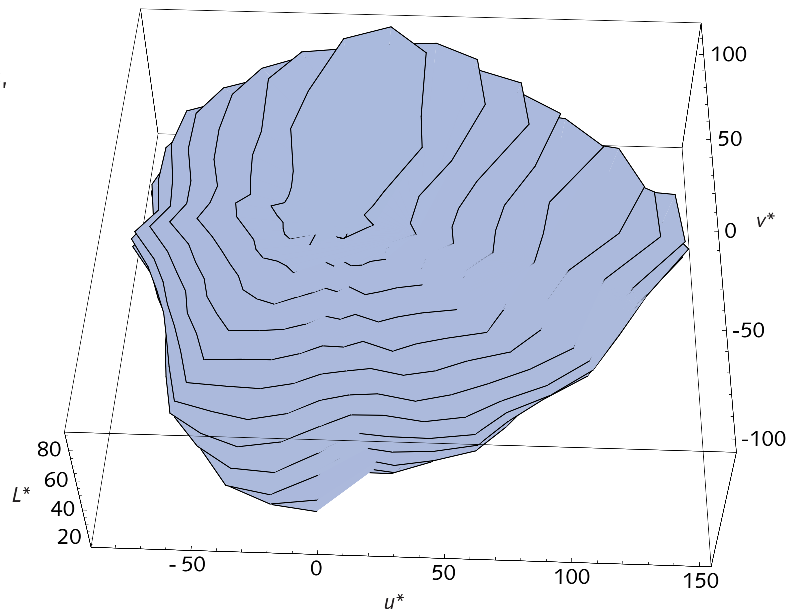

linear-light domain, prior to gamma correction. (Some proponents of digital cinema recommend capture with the matrix "switched off." I will address that view later.)

Colour cameras deliver three components, and obviously those three components are interpreted by the display as representing primaries of known chromaticities. What is their relationship with the camera signals? Does the camera have primaries? The answers to these questions are not unanimously agreed upon by colour scientists: "Experts disagree!" In the following sections, I will give my interpretation.

Figure 5 **Coordinates of ColorChecker patches** are graphed on the CIE [*x*, *y*] chromaticity diagram. The horseshoe encloses all colors; the triangle encloses the colors that can be represented in video (BT.709) and in desktop computing (sRGB).

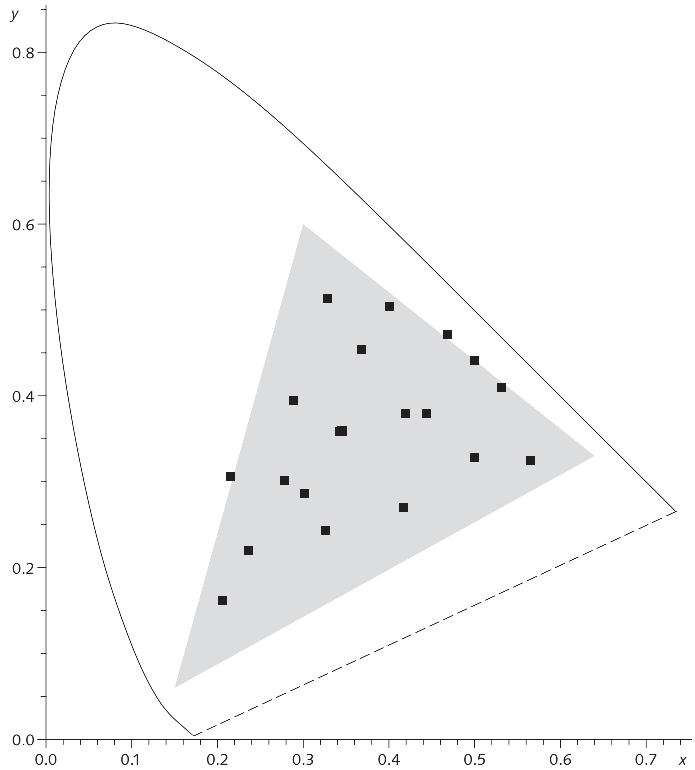

Interpretation of raw camera *RGB*

If you pay no attention to colour science, and simply connect a camera's output signals to a monitor – or to the front end of a postproduction chain – you will get colours. However, the colours displayed will not, except in unusual circumstances, be very close to those in the scene. Figure 5 shows the chromaticities of the 24 patches of the GretagMacbeth ColorChecker as measured by a colourimeter. When imaged by a typical digital camera, and the uncorrected *R*, *G*, and *B* values are treated as BT.709 values, the chromaticities of Figure 6 result. The most obvious deficiency is that the uncorrected device values exhibit a loss of colour saturation. The loss of saturation occurs mainly because negative sensitivities at certain wavelengths would be required to implement an "ideal" sensor for the BT.709 display primaries.

Negative lobes

The necessity of "negative lobes" is explained in the passages on pages 240 through 243 of *DVAI* [1], and in the accompanying 6-frame "storyboard" set of graphs and captions on pages 244 through 249. I'll summarize the argument in the remainder of this section.

If capture was performed with the CIE $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$ colour matching functions (CMFs), then all colours would be captured, and all colours would be represented in nonnegative *XYZ* values. However, direct display of these *XYZ* values would require negative power at certain wavelengths at the display. In other words, direct display would require nonphysical (nonrealizable) SPDs at the display.

For physical (realizable) SPDs at the display – say, for display using BT.709 primaries – it is relatively straightforward to work out the CMFs required to accurately capture suitable signals. Figure 4 (on page 6) shows the ideal CMFs required for colour signals to be captured for BT.709 display. The transformed CMFs required at the camera (in this case, BT.709 CMFs) inevitably have negative lobes – obviously a problem for a 3-channel camera!

The spectral responsivities of Figure 7 could be implemented by a 6-channel camera having a set of 3 channels sensitive to the positive lobes of each of the three CMFs augmented by a set of 3 channels sensitive to the negative lobes of

Figure 6 **Uncorrected device RGB values** of a typical digital camera are graphed here as if the camera produced *RGB* tristimulus values corresponding to the sRGB primaries. The most obvious problem is that the patches are apparently desaturated. Signal processing can be used to bring these values into closer agreement with the values obtained using the CIE Standard Observer.

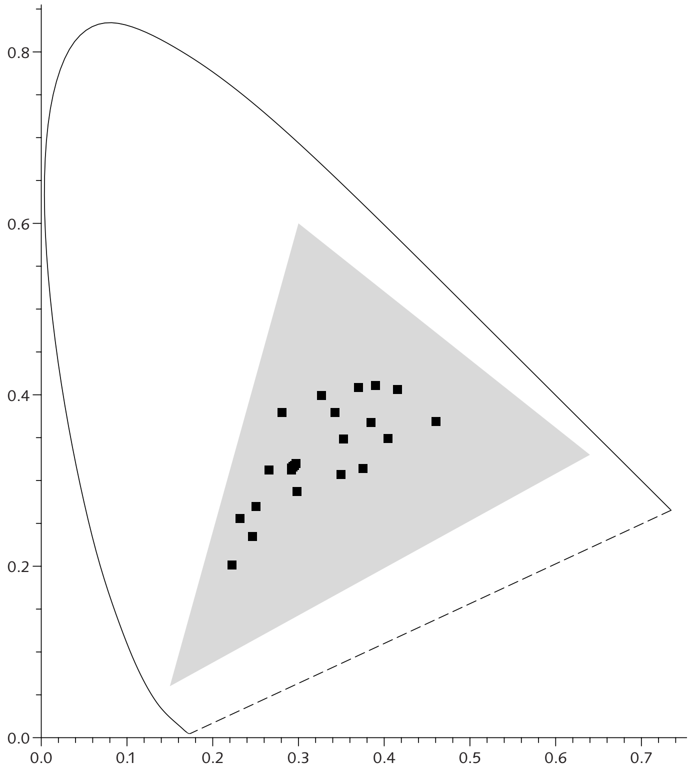

Figure 7 **CMFs for Rec. 709** are the theoretically correct analysis functions to acquire *RGB* components for display using Rec. 709 primaries. Owing to their negative lobes, they are not directly realizable in a camera or a scanner, but they can be realized through use of the the CIE *XYZ* color matching functions (or any nondegenerate linear transformation of them) followed by signal processing involving a 3×3 matrix transform.

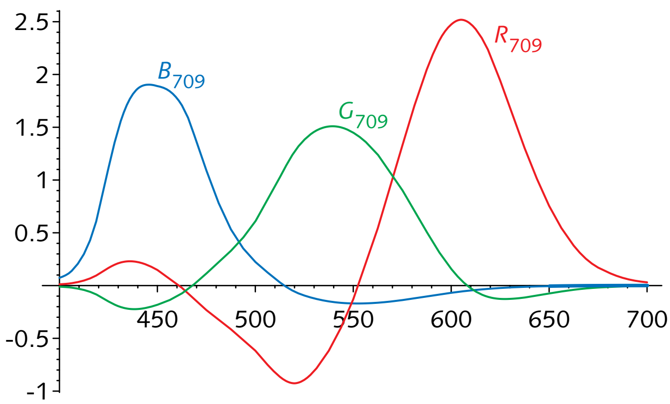

each of the three CMFs. Signal components of the negative-going channels (or channel) could then be electrically negated and summed with the corresponding positive-going signals.

For BT.709's CMFs, the green channel's two negative lobes and the blue channel's single negative lobe are quite low in amplitude. An engineer would be tempted to ignore these, and to ignore the small secondary positive lobe of red. That approach would lead to a four-channel camera;[7] a fraction of the fourth channel's signal would be subtracted from the other 3.

However, four channels aren't necessary. Just three channels implement $\bar{x}\bar{y}\bar{z}$ capture can be

7 I speculate that the Sony RGB+E camera uses three SRFs similar to the dominant positive lobes of Figure 7, and a fourth SRF comparable to the inversion of the large negative lobe of R_{709} (which lies in the cyan region of the spectrum).

Figure 8 **Typical digital camera SRFs** (spectral responsivity functions) are graphed here. The red, green, and blue channels are graphed in the corresponding colors. Because these responses differ from the CIE standard observer, the native device values cannot be sensibly interpreted as *XYZ;* because the SRFs differ from the BT.709 CMFs, the native device values cannot be sensibly interpreted as sRGB. However, with the application of a suitably optimized linear 3×3 matrix, reasonably accurate *XYZ* or sRGB color information can be obtained.

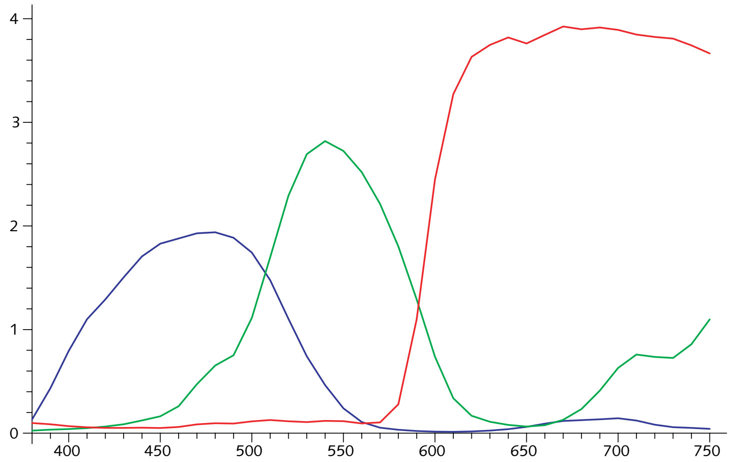

used. Instead of negating three channels and summing, or subtracting some fraction of a fourth channel, a full 3×3 matrix (with some negative coefficients) is used.

If ideal $\overline{xyz}$ capture was performed, producing *XYZ* signals, this 3×3 matrix would be required to encode into BT.709 *RGB* signals:

$$\begin{bmatrix} R_{709} \\ G_{709} \\ B_{709} \end{bmatrix} = \begin{bmatrix} 3.240479 & -1.537150 & -0.498535 \\ -0.969256 & 1.875992 & 0.041556 \\ 0.055648 & -0.204043 & 1.057311 \end{bmatrix} \bullet \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

Eq 1

Processing through a matrix such as that of Equation 1 has noise implications. The top left coefficient of that matrix, about 3.24, causes 1 mV (or 1 code value) of noise in the *X* channel to be amplified into 3.24 mV (or 3 code values) in the resulting *R* signal. The large overlap between the $\bar{x}$ and $\bar{y}$ sensitivities produces the large departure from an identity matrix. From a noise perspective, the optimum 3×3 matrix would be the identity matrix.

Practical cameras do not have $\overline{xyz}$ sensitivities. Instead, camera designers tune their colour separation filters to depart from $\overline{xyz}$ sensitivities and tune their matrices for a balance between low noise, acceptable metamerism, and reasonably accurate colour. Users are destined to live with the camera metamerism that results from failing to adhere to the Maxwell-Ives criterion.

Optimum 3×3 matrices

Construction of optimum 3×3 matrices is a combination of science, craft, and perhaps even art. At its simplest, you start with a coloured optical stimulus such as the ColorChecker chart.[8] You measure the patches with a colour measuring instrument, and use the parameters of the intended target colourspace (e.g., BT.709) to compute a set of idealized target *RGB* values. Then you use your camera to capture the stimulus and obtain actual, native device values. Finally, you can construct a colour transform that maps the native device values to the target *RGB* values according to some optimization criteria. For reasons that I'll detail later, in my view the best transform is a 3×3 "linear matrix."

The simplest form of determining an optimum 3×3 matrix involves least-squares techniques. Given a matrix ***D*** whose columns contain sets of device *RGB* triples, and a matrix ***R*** containing the corresponding ideal target *RGB* triples, a 3×3 matrix ***M*** maps from ***D*** to ***R***:

$$\boldsymbol{R} = \boldsymbol{M} \cdot \boldsymbol{D} \qquad \text{Eq 2}$$

The optimum matrix ***M*** is found by solving Equation 2, either directly (the preferred approach), or by computing the matrix pseudoinverse of ***R*** then computing the matrix product

8 You can call the collection of optical stimuli (or its synthetic equivalent) a *training set.*

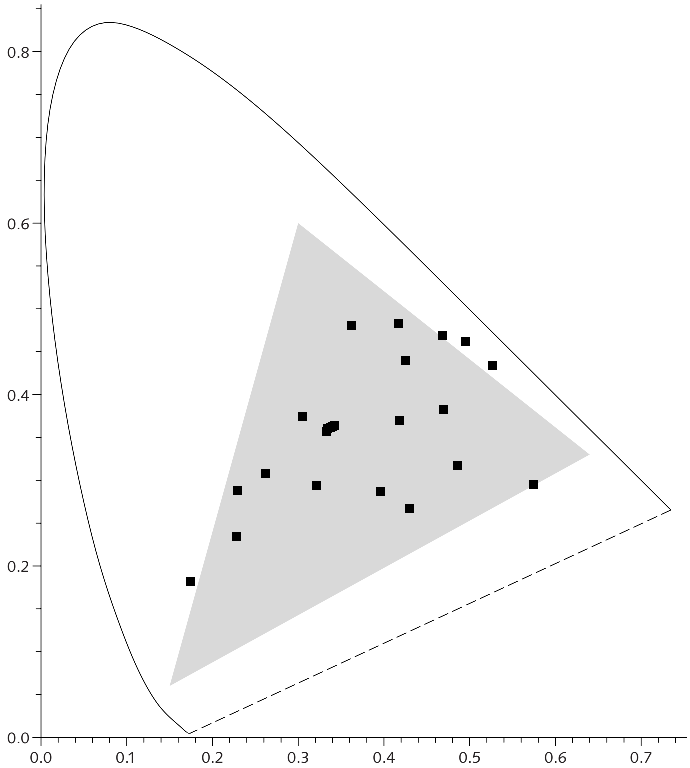

Figure 9 **Corrected device values,** after mapping through an optimum 3×3 matrix, are graphed here. The chromaticity values are reasonably close to those of Figure 5.

(by premultiplication, that is, left-multiplication) with ***D***:

$$\boldsymbol{M} = \boldsymbol{D} \cdot \boldsymbol{R}^{+} \qquad \text{Eq 3}$$

Figure 8 shows spectral responsivity functions (SRFs) of a typical digital camera. For that camera, this matrix results:

$$\begin{pmatrix} \hat{R} \\ \hat{G} \\ \hat{B} \end{pmatrix} = \begin{pmatrix} 0.7659 & 0.7052 & -0.4990 \\ -0.3140 & 1.3283 & -0.1367 \\ 0.0609 & -0.4739 & 1.0326 \end{pmatrix} \bullet \begin{pmatrix} R_{\mathrm{TCS230}} \\ G_{\mathrm{TCS230}} \\ B_{\mathrm{TCS230}} \end{pmatrix} \qquad \text{Eq 4}$$

Notice the large off-diagonal terms – having magnitudes up to 0.7 – and fairly large negative terms. Figure 9 shows the result of mapping the ColorChecker patches through the optimum matrix of Equation 4. Evidently the ColorChecker patches are mapped to chromaticity coordinates reasonably close to their ideal coordinates as shown in Figure 5. The optimum matrix yields an average error for this particular camera of about 5 ΔE_{ab}. (Keep in mind that these 2-D representations do not portray the mapping of luminance levels.)

Refinements

I have outlined the pseudoinverse technique. Many refinements of this technique can be, and are, used in computing optimum linear 3×3 matrices. I'll briefly outline a few refinements and alternate approaches:

- The *principal eigenvectors* (PE) method, also known as *truncated SVD*, involves discarding from the "training set" those elements that are determined, from the mathematical procedure, not to contribute significantly to the estimated matrix coefficients. Such samples are discarded because they are likely to contribute noise.
- The least-squares weighting can be weighted according to colours for which it is especially important to maintain accuracy. For example, the least-squares solution can be weighted to emphasize accurate mapping of skin tones. (Other colours will necessarily suffer.)
- It may be important that the grayscale maps correctly. Correct mapping of grays includes white, of course. The refinement, detailed by Finlayson and Drew [11], is called *white point preserving least-squares* (WPPLS).
- I have described using a real-life optical stimulus – the ColorChecker. If actual spectral

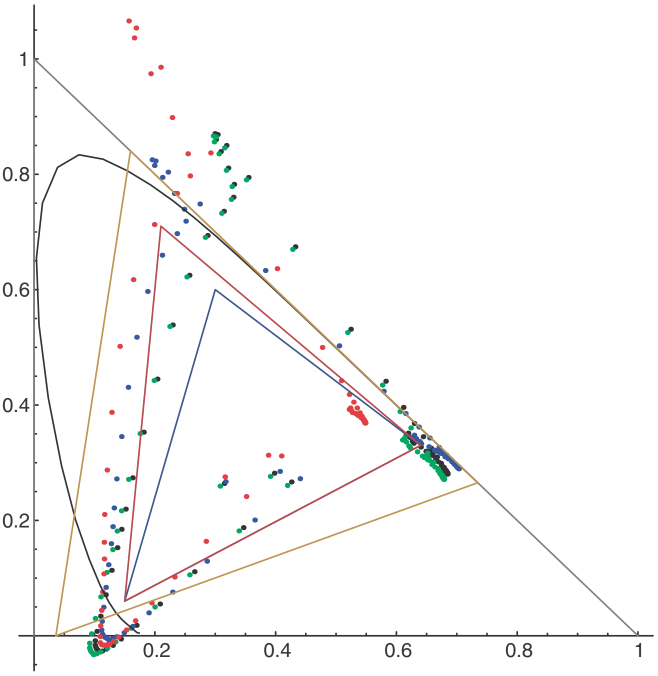

Capture color analysis gamuts for a Canon 20D digital camera (LS matrix-black dots, WPPLS matrix-green dots, RGB error minimization matrix-blue dots, DNG D65 matrix-red dots)

Figure 10 **Capture color analysis gamut** is illustrated in this sketch taken from Jack Holm's paper cited in the text.

responsivity data (SRFs) of the camera is available, the calculations can be done synthetically.

- If actual SRFs are available, a synthetic analysis can be performed on monochromatic spectral stimuli instead of the ColorChecker. [12] Think of this approach as using 31 test stimuli (31 "patches"), where each stimulus is concentrated at a single wavelength. This approach is mathematically equivalent to finding the 3×3 matrix that best matches, in a least-squares sense, the ideal CMFs for the intended image encoding primaries. (For example, if targeting BT.709, the technique finds the linear combination of native device SRFs that best matches Figure 7.) Some researchers argue that using monochromatic stimuli ought to give better performance: According to their view, optimization performed at the spectral boundary ought to better constrain colour mappings within the entirety of colourspace. Other researchers argue that spectral stimuli will never be encountered in actual use of the camera, and that it is more important to optimize for realistic stimuli. (I tend toward the latter view.) In the limit, *in situ* scene-dependent illumination SPDs and spectral reflectances could be used.
- The procedure that I have described implies that the error being minimized is what you might denote ΔXYZ or ΔRGB, in linear-light space. It may be more appropriate to minimize a more perceptual error metric such as ΔE_{ab} (that is, an error measured in CIELAB). Delta-E is a nonlinear function of XYZ (or RGB), so a nonlinear optimizer is necessary. The Nelder-Mead technique (implemented in Excel's *Solver,* MATLAB's *fminsearch*, and Mathematica's *NMinimize*) could be used. [13]
- Finally, error minimization could be performed in a colour appearance space such as CIECAM02 *Jab.*

Wide-gamut capture

I can now summarize my conclusions concerning colour capture:

- *Any* set of SRFs captures all colours. A camera *per se* does not limit capture gamut: A three-component camera potentially has unlimited gamut.
- Camera metamerism will be present to the extent that the SRFs depart from the CIE CMFs.
- Metamerism should not be confounded with capture gamut.

- For reasonably well-controlled illuminant spectra (as is typical of professional image capture), and absent any pathological spectral reflectances in the scene, metamerism is not a serious problem.
- Colour mapping accuracy is dependent upon the camera SRFs, and is influenced by illumination spectra and spectral reflectance of scene elements.

"Optical" primaries

A 3×3 linear matrix that transforms from *XYZ* to an *RGB* primary system, say RGB_1, can be constructed as a function of the chromaticities and the primaries of the *RGB* system. The procedure and the math are described on page 250 of *DVAI* [1]. As an extension of this technique, the 3×3 linear matrix that transforms from one primary system (say RGB_1) to another (say RGB_2) can be constructed by inverting the RGB_1-from-*XYZ* matrix described above, and concatenating that matrix with the matrix that transforms from *XYZ* to RGB_2 (using the chromaticities and the primaries of RGB_2).

If we have constructed an optimum 3×3 matrix that maps from native device *RGB* to *a* standard interchange *RGB* (such as BT.709), then it is trivial to construct the transformation to *XYZ*.

The 3×3 construction technique that I have described can be reversed: We can take the optimum 3×3 matrix, expressed in *XYZ*-from-*RGB* form, and "deconstruct" it to extract the chromaticities of the primaries and white. That procedure will obtain the primaries and white that would (whether realizable or not) produce colour correctly without using a 3×3 matrix. I call these the "optical" (or "native") primaries[9] of the camera. These primaries can be plotted on a chromaticity diagram.

When this procedure is undertaken for SRFs and 3×3 matrices of actual cameras, the "optical" primary chromaticities generally lie somewhat outside the spectral locus.

9 What I call the "optical" primaries are dependent upon the assumptions of the optimization procedure (perhaps including estimated scene statistics), and a function of the assumed illuminant SPD.

Future directions

It remains a research problem to investigate capture noise as a function of the combination of colour separation filters and 3×3 matrices. I believe that capture noise is a poorly understood constraint on selection of camera separation filters.

It would be useful in the long term to investigate the statistics of illuminants and scene elements, because separation filters and 3×3 matrices depend upon these statistics.

I expect the positions of the optical primaries on a chromaticity diagram to give insight into colour and noise performance of cameras. ■

References

1 Poynton, Charles (2003), *Digital video and HDTV algorithms and interfaces* (San Francisco: Morgan Kaufmann).

2 Cohen, Jozef B., and William E. Kappauf (1985), "Color mixture and fundamental metamers: Theory, algebra, geometry, application," in *Am. Journal of Psychology,* **98** (2): 171–259 (Summer).

3 Ejaz, Tahseen, et al. (2006), "Development of a Camera System for the Acquisition of High-Fidelity Colors," in *IEICE Trans. Electron.* **E89-C** (10): 1441–1447.

4 MacAdam, David L. (1993), "Maximum Visual Efficiency of Colored Materials," in *J. Opt. Soc. Am.* **15**: 361–367 (Nov. 1935); reprinted in MacAdam, David L. (ed.), *Selected Papers on Colorimetry – Fundamentals* (Bellingham, Wash.: SPIE Press).

5 Martínez-Verdú, Francisco, et al. (2006), "Calculation of the Optimal Colors of Linear Input Devices," in *Proceedings of CGIV 2006* (IS&T, Leeds, U.K.), 345–349 (June).

6 Stiles, W.S. and G.W. Wyszecki (1962), "Counting metameric object colors," in *J. Opt. Soc. Am.* **52** (3): 313–328 (Mar.).

7 Stiles, W.S., et al. (1977), "Counting metameric object-color stimuli using frequency-limited spectral reflectance functions," in *J. Opt. Soc. Am.* **67** (6): 779–784 (June).

8 Pointer, M.R. (1980), "The gamut of real surface colors," in *Color Research and Application* **5** (3): 145–155 (Fall).

9 Centen, Peter (2004), "How Wide Gamut Is A Broadcast Camera," in *Proc. 146th SMPTE Tech. Conf.* (Pasadena, Calif., Oct.).

10 www.cis.rit.edu/mcsl

11 Finlayson, G.D. and M.S. Drew (1997), "White-point preserving color correction," in *Proc. IS&T/SID 5th Color Imaging Conference* (Scottsdale, Ariz.): 258–261 (Nov.).

12 Holm, Jack (2006), "Capture Color Analysis Gamuts," in *Proc. IS&T/SID 14th Color Imaging Conference* (Scottsdale, Ariz.): 108–113 (Nov.).

13 Funt, Brian, et al. (2004), "Optimal Linear *RGB*-to-*XYZ* Mapping for Color Display Calibration," in *Proc. IS&T/SID 12th Color Imaging Conference* (Scottsdale, Ariz.): 223–227 (Nov.). This paper concerns displays, but the technique is applicable to cameras.