


Sensing color with the TAOS TCS230

The TAOS TCS230 is a small, highly integrated color sensing device packaged in a clear plastic 8-pin SOIC. It reports, as analog frequency, the amount of shortwave (blue), mediumwave (green), longwave (red), and wideband (white) optical power incident onto the device. It can be used in a variety of color sensing applications. Details of the device can be found in its datasheet. This white paper details the concepts and calculations involved in color sensing using the TCS230.

www.gretagmacbeth.com

We will use the ColorChecker chart as an optical stimulus to work through a numerical example of color sensing. The chart, depicted in Figure 1, is manufactured and distributed by GretagMacbeth. The chart measures approximately 13 inches by 9 inches (330 mm by 230 mm); it contains 24 colored patches arranged in a 6 by 4 array. Figures 2 through 5 overleaf show the spectral reflectance of the patches in each of the four rows of the chart – that is, the fraction of incident light that is reflected (with respect to an ideal diffuse reflector), as a function of wavelength from 350 nm to 750 nm.

Figure 1 **The ColorChecker** contains 18 colored patches and a 6-step gray series. Whether the colors here display or print correctly depends upon color management at your display and at your printer.






Figure 2 ColorChecker spectra, top row.




Figure 3 ColorChecker spectra, second row.




Figure 4 ColorChecker spectra, third row.




Figure 5 ColorChecker spectra, bottom row (neutral series)




Figure 6 **Cone sensitivities** of cone photoreceptors are shown. The shortwave-sensitive photoreceptors are much less sensitive than the other two types. The responses of the mediumwave and longwave photoreceptors have a great deal of overlap. Vision is not sensitive to the precise wavelength of the stimulus: What matters is optical power integrated under each response curve.

Introduction to color vision

Photoreceptor cells called *cones* in the retina are responsible for human color vision. There are three types of cone cells, sensitive to longwave, mediumwave, and shortwave radiation within the electromagnetic spectrum between about 400 nm and 700 nm. Because the cone sensitivities are very roughly in the parts of the spectrum that appear red, green, and blue, color scientists denote the cell types as ρ , γ , and β , the Greek letters for r, g, and b. (To denote the sensors R , G , and B would wrongly suggest a closer correspondence.) Estimates of the spectral response of the cone types are graphed in Figure 6 above.

Sometimes the designations S , M , and L – for shortwave, medium-wave, and longwave – are used.

Light in the physical world can be characterized by spectral power distributions (SPDs). Colored objects can be characterized by spectral reflectance curves, such as those of the ColorChecker. However, vision is insensitive to the exact wavelength of a stimulus: According to the modern theory of color science, all that matters is the integral of optical power underneath each response curve. That there are exactly three types of cone cells leads to the property of *trichromaticity*: Three components are necessary and sufficient to characterize color. Some people might use the phrase “color as sensed by the eye,” but I consider that qualifier to be redundant at best, and misleading at worst: Color is *defined* by vision, so there is no need to use the qualifying phrase “as sensed by the eye,” or to use the adjective *visible* when referring to color.

Overview of CIE Colorimetry

The spectral responses of the cone cells that I graphed in Figure 6 were unavailable to researchers in the 1920s. Researchers at the time used psychophysical experiments, such as the famous color matching experiment, to tease out the data. The CIE is the international body responsible for color standards. In 1931, that organization adopted the *color matching functions* denoted $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$, graphed in Figure 7.




Figure 7 **CIE 1931, 2° color-matching functions.** A camera with 3 sensors must have these spectral response curves, or linear combinations of them, in order to capture all colors. However, practical considerations make this difficult. These analysis functions are *not* comparable to spectral power distributions!

A description of the application of color science to color imaging is contained in this book:

Poynton, Charles, *Digital Video and HDTV Algorithms and Interfaces* (San Francisco: Morgan Kaufmann, 2003).

Weighting a physical SPD under each of these three curves (that is, forming the wavelength-by-wavelength product), and summing the results, forms a triple of three numbers, denoted X , Y , and Z . In continuous mathematics, three integrals need to be computed; in discrete math, a matrix product is sufficient. The X , Y , and Z *tristimulus values* characterize color. They are linear-light quantities, proportional to optical power, that incorporate the wavelength sensitivity of human vision. The Y value is *luminance*, which is ordinarily expressed in units of candela per meter squared ($\text{cd} \cdot \text{m}^{-2}$). If you are measuring reflectance, the reflected tristimulus values depend upon the spectral characteristics of the illuminant, and their amplitudes scale with the power of the illumination. *Relative luminance* is the ratio of reflected luminance to the luminance of the illumination; it is also known as the *luminance factor*.

The $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$ color matching functions of Figure 7 are obviously different from the ρ , γ , and β spectral sensitivities of Figure 6. The two sets are related by a 3×3 linear transform. The functions of Figure 6 were not known in 1931 when the CIE adopted the standard color matching functions. The “bump” in the $\bar{x}(\lambda)$ function is a consequence of a normalization process; it is clear from Figure 6 that the bump is not a feature of the cone response.

Luminance is proportional to intensity, but is weighted by the spectral response of vision's lightness sensation. The perceptual response to luminance is complex, but it can be approximated by a power func-




Figure 8 **SPDs of various illuminants** are graphed here. Illuminant A, shown in orange, is representative of tungsten light sources; it is deficient in shortwave power, and may cause errors in sensing blue colors. The blue line graphs the SPD of a Nichia white LED. There is a peak in the blue portion of the spectrum: Uncorrected, the sensor would report excessive blue values. The other four lines represent CIE standard illuminants C, D₅₀, D₅₅, and D₆₅.

tion. The CIE has standardized the computation of *lightness*, denoted L^* , as approximately the 0.4-power of relative luminance. Because of the subtle relationships among intensity, luminance, and lightness, you should take care to use these terms correctly.

$$x = \frac{X}{X+Y+Z}$$

$$y = \frac{Y}{X+Y+Z}$$

Eq 1 Chromaticity coordinates

In many applications, tristimulus signals (including luminance) scale with the illumination, and are otherwise uninteresting in themselves. What is more interesting is the ratios among them, which characterize color disregarding luminance. The CIE has standardized the projective transformation of Equation 1, in the margin, to transform $[X, Y, Z]$ values into a pair of $[x, y]$ *chromaticity coordinates* that represent color disregarding luminance. These coordinates are suitable for plotting in two dimensions on a *chromaticity diagram*.

Illumination

A nonemissive object must be illuminated in order to be visible. The SPD reflected from an illuminated object is the wavelength-by-wavelength product of the illuminant's SPD and the spectral reflectance of the object. Before light reaches the eye, the interaction among light sources and materials takes place in the spectral domain, not in the domain of trichromaticity. To accurately model these interactions requires spectral computations. When applying the TCS230, attention must be paid to the spectral content of the illumination and to potential interaction between the illumination and the samples to be sensed. Generally, the less spiky the spectra, the better. Figure 8 graphs several illuminants.

Your application may involve sensing color, in which case the preceding description applies. However, some applications of the TCS230 involve not so much estimating color as seen by the eye but rather sensing physical parameters associated with optical power in the visible range. In such applications, to approximate the visual response may not be the best approach: It may be more effective to take a more direct approach to estimating the parameters of the underlying physical process.

The Color Checker

Equipped with knowledge of how spectra are related to colors, the plotting of chromaticity coordinates, and the dependence of colors upon illumination, we can return to the ColorChecker. GretagMacbeth doesn't publish or guarantee the spectral composition of the patches of the ColorChecker. However, nominal CIE $[X, Y, Z]$ values are published. The patches in the bottom row of the ColorChecker contain neutral colors; the numeric notations in the legends of Figure 5 reflect one tenth of the lightness (L^*) values of those patches.

The spectra graphed on pages 2 and 3 represent the physical wavelength-by-wavelength reflectance of the patches. These spectral reflectances have been measured by color measurement instrument called a *spectrophotometer*. If you had access to a light source having perfectly even distribution of power across the visible spectrum, then the reflectance curves graphed here could simply be scaled to represent the reflectance in your application. Practical light sources do not have perfectly even spectral distributions, so compensation is necessary: You must compute the wavelength-by-wavelength product of the illuminant's SPD with the spectral reflectance of the chart.

We will first calculate the CIE $[X, Y, Z]$ values from the chart. (These values should agree with the figures provided by Gretag.) Then we will calculate the $[R, G, B]$ values that will be detected by a TCS230.

To calculate CIE $[X, Y, Z]$, we take the 31×3 matrix representing the color matching functions (CMFs) of the CIE Standard Observer, and perform a matrix product with 31 spectral response values as corrected for illumination. This produces the $[X, Y, Z]$ tristimulus values. When chromaticity coordinates $[x, y]$ are computed from $[X, Y, Z]$ through the projective transform in Equation 1, then plotted, the chromaticity diagram in Figure 9 results. The horseshoe-shaped figure, closed at the bottom, contains all colors: Every non-negative spectral distribution produces an $[x, y]$ pair that plots within this region. The lightly-shaded triangle shows the region containing all colors that can be produced by an additive RGB system using sRGB (Rec. 709) primary colors. This region typifies video and desktop computing (sRGB). The points plotted in Figure 9 are the colors of the ColorChecker. White and gray values are clustered near the center of the chart.




Figure 9 **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 (Rec. 709) and in desktop computing (sRGB).

The TCS230

Figure 10 shows the responses of the four channels of the TCS230. The black curve shows the response of the unfiltered sensor elements. The red, green, and blue curves show the responses of the longwave-sensitive, mediumwave-sensitive, and shortwave-sensitive elements respectively.

As I mentioned on page 5, the CIE model of color vision involves integrating an SPD under the $X(\lambda)$, $Y(\lambda)$, and $Z(\lambda)$ color matching functions (graphed in Figure 7), producing X , Y , and Z values. To use the TCS230 to estimate color we perform an analogous calculation, but using the TCS230 sensitivity functions instead of the CIE CMFs: We integrate the SPD under the TCS230's sensitivity curves, and produce R , G , and B values. The device R , G , and B values will depend upon several factors: the spectral content of the illuminant, the spectral reflectance of the sample, the spectral attenuation of any intervening optical components (such as the lens), and finally, the spectral response functions of the TCS230. The various spectral phenomena are modelled by computing wavelength-by-wavelength products.

Owing to the fact that the TCS230 is sensitive to infrared light (having wavelengths above 700 nm), and the fact that most light sources




Figure 10 **TCS230 spectral sensitivities** are graphed here. The red, green, and blue channels are graphed in the corresponding colors; the gray line reflects the sensitivity of the clear (unfiltered) channel. Because these responses are different from the CIE standard observer, the values reported by the TCS230 are not colorimetric. However, suitable signal processing yields color information that is sufficiently accurate for many industrial applications.

produce power in the infrared region, typical applications include an IR cut filter in front of the TCS230. Figure 11 overleaf shows the response of a typical IR cut filter.

Continuing our modelling of the ColorChecker, we illuminate the ColorChecker with the CIE D_{65} illuminant, integrate the resulting spectral reflectances under the TCS230 sensitivity curves, and finally transform to CIE $[x, y]$ coordinates. The relative luminance values obtained through this process are fairly accurate; however, the chromaticity coordinates are not very accurate. Figure 12 (on page 11) graphs the CIE chromaticities of the uncorrected R , G , and B values. The results differ from the coordinates of the ColorChecker graphed in Figure 9.

The reason for the disagreement is that the TCS230's sensitivity functions differ quite substantially from the color matching functions that are appropriate for sRGB. Even if the TCS230 sensitivities were in close agreement with sRGB, the effect of the spectral power distribution of the illuminant and the spectral effect of intervening optical components would cause some divergence.

To form a more accurate estimate of color requires processing the raw TCS230 R , G , and B values through a linear 3×3 matrix whose coefficients are optimized with respect to the spectrum of the illuminant, the spectral response of intervening optical components, and the




Figure 11 **An IR cut filter** is necessary in most sensor systems. Most silicon sensors are sensitive to IR, most light sources emit substantial amounts of power in the IR region, and many colored objects reflect in the IR range. If IR response were left unattenuated, the sensor would report excessively high red components. This graph reflects the IR filter incorporated into the Argus lens of the TCS230 evaluation module.

response curves of the TCS230. The data processing operation can be represented in matrix form as follows:

$$\mathbf{x} = \mathbf{M} \bullet \mathbf{t} \quad \text{Eq 2}$$

I conform to the usual mathematical convention of representing a vector by a lowercase bold italic letter and a matrix by an uppercase bold italic letter.

The symbol \mathbf{t} represents a three-element vector containing the device values captured from a color patch. \mathbf{M} represents the 3×3 color correction matrix that we will apply to these values through matrix multiplication, denoted by the \bullet symbol. The symbol \mathbf{x} represents the resulting vector of estimated $[X, Y, Z]$ values.

We can use matrix notation to symbolize processing a set of three color patches at once, by arranging the three sets of device values into successive columns of a 3×3 matrix \mathbf{T} . Successive rows of \mathbf{T} contain red, green, and blue data respectively. Upon matrix multiplication by \mathbf{M} , the columns of the resulting matrix \mathbf{X} contain XYZ values of the successive samples; the rows of \mathbf{X} contain X, Y, and Z values respectively. One equation expresses the mapping of three patches at once:

$$\mathbf{X} = \mathbf{M} \bullet \mathbf{T} \quad \text{Eq 3}$$

Given a matrix \mathbf{T} whose columns contain three sets of device samples, and a matrix \mathbf{X} containing the corresponding set of three ideal XYZ triples, there is a unique matrix \mathbf{M} that maps from \mathbf{T} to \mathbf{X} . It is found




Figure 12 **Uncorrected TCS230 values**, illuminated by CIE Illuminant D_{65} , are graphed here as if the TCS230 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.

by computing the matrix inverse of T , then computing the matrix product (by premultiplication) with X :

$$M = X \cdot T^{-1} \quad \text{Eq 4}$$

Matrix multiplication is non-commutative: $A \cdot B$ is generally unequal to $B \cdot A$. *Premultiplication* by X means that X is on the left.

The resulting 3×3 color correction matrix M exactly maps the each of the chosen three sets of device values to the corresponding set of tristimulus values. It is not necessary to invert matrices at the time of sensing! The matrix M can be computed in advance, based upon the samples that are expected to be presented to the sensor in the intended application. To process three device values upon sensing a sample, all that is necessary is computation of the matrix product of Equation 3.

Here is the transform to CIE $[X, Y, Z]$ for the ColorChecker's red, green, and blue patches, illuminated by CIE D_{65} and sensed by the TCS230:

$$\begin{pmatrix} \hat{X} \\ \hat{Y} \\ \hat{Z} \end{pmatrix} = \begin{pmatrix} 0.3568 & 0.2889 & 0.2132 \\ 0.1173 & 0.9000 & -0.0611 \\ 0.0095 & -0.5294 & 1.6196 \end{pmatrix} \cdot \begin{pmatrix} R_{TCS230} \\ G_{TCS230} \\ B_{TCS230} \end{pmatrix} \quad \text{Eq 5}$$

The “hats” on X , Y , and Z indicate that these quantities are estimates. The matrix in Equation 5 was optimized for three particular sample values (patches). The approach is suitable for applications where the TCS230 is used to distinguish among a set of three colors. However, exact mapping is guaranteed for just those three samples. The mapping from a set of device values outside that set is likely to be far from optimal. Attempting to optimize the mapping for more than three samples, to a set of more than three corresponding ideal tristimulus values, leads to an overdetermined system. If your application involves more than three stimuli, further work is necessary.

A color correction matrix that produces good results across more than three samples can be computed through a numerical optimization procedure. When this is done, no particular sample is likely to map exactly to its ideal tristimulus set, but a linear matrix can be constructed that minimizes the error across a range of samples (where the error is measured in a least-squares sense). The color correction operation is still accomplished exactly as in Equation 2.

The mathematical details of the pseudoinverse are explained in Strang, Gilbert, *Introduction to Linear Algebra*, Second Edition (Boston: Wellesley-Cambridge, 1998). For a highly technical description of the construction of color transforms, see Sharma, Guarav, Ed., *Digital Color Imaging Handbook* (Boca Raton, FL, U.S.A.: CRC Press, 2003).

To describe the mapping of more than three patches using just one equation, matrix T is extended to have more than three columns. In this example, we will use the 24 patches of the ColorChecker chart, so T is a 3×24 matrix and X is a 3×24 matrix. The numerical optimization procedure takes exactly the form of Equation 4; however, the system is overdetermined, and the exact inverse of T does not exist. The best least-squares mapping is obtained by computing the *pseudoinverse* of T instead of the inverse. Pseudoinverse is related to the *singular value decomposition* (SVD). Systems for doing mathematics using computers – such as Mathematica and MATLAB – have built-in provisions for computing the pseudoinverse of a matrix.

When the pseudoinverse is formed and premultiplied by the ideal tristimulus values – that is, multiplied on the left of the pseudoinverse – the optimum transform for the ColorChecker illuminated by CIE D₆₅ and sensed by the TCS230 is determined to be this:

$$\begin{pmatrix} \hat{X} \\ \hat{Y} \\ \hat{Z} \end{pmatrix} = \begin{pmatrix} 0.4287 & 0.4489 & 0.0493 \\ 0.1450 & 0.9623 & -0.1167 \\ 0.0539 & -0.4059 & 1.4191 \end{pmatrix} \bullet \begin{pmatrix} R_{TCS230} \\ G_{TCS230} \\ B_{TCS230} \end{pmatrix} \quad \text{Eq 6}$$

Euclidean distance is the square root of the sum of the squares of a set of values. In engineering, this is the *root-mean-square* (RMS) value. For color differences, we deal in three dimensions. In the two-dimensional case, this computes the length of the hypotenuse of a right triangle. (Upon receiving his diploma, the Scarecrow in *The Wizard of Oz* offers Dorothy a garbled description of this calculation; he has gained self-confidence, but not mathematical knowledge!)

ColorChecker patches processed through this optimized transform produce the chromaticity coordinates that are plotted in Figure 13.

The usual way to describe the visual magnitude of color differences is to use the *delta-E* scale defined by the CIE. This scale is defined by Euclidean distances between CIE $L^*a^*b^*$ triples, where $L^*a^*b^*$ values – sometimes denoted *CIELAB*, or just *LAB* – are obtained by a nonlinear transformation of CIE *XYZ*. Transforms among *XYZ*, *Yxy*, $L^*a^*b^*$, *RGB*, and many other color systems are detailed in Poynton's book, cited in the margin of page 5.




Figure 13 **Corrected TCS230 values**, illuminated by CIE Illuminant D_{65} , are graphed here. The chromaticity values are reasonably close to those of Figure 9.

When color-corrected TCS230 values are transformed to $L^*a^*b^*$ coordinates through the optimized matrix of Equation 6, then compared to the published $L^*a^*b^*$ values of the ColorChecker (as measured by instrumentation), the average magnitude of the error is about 5.5 delta- E units. In high-quality image reproduction, a delta- E value of unity is taken to lie approximately on the threshold of perceptibility, and a value of 2 is taken to be acceptable. A color measurement instrument is expected to have an error of a fraction of a delta- E . In TCS230 applications, we are not sensing images. An average error of 5.5 is quite respectable for a low-cost sensor device, and is comparable to the color error present in image data sensed by consumer digital still cameras.

You may wish to transform TCS230 color estimates to sRGB values suitable for use in desktop computing or computer graphics. You can transform $[X, Y, Z]$ values to linear-light R , G , and B tristimulus values having the appropriate primary and white reference chromaticities of the sRGB standard by using this linear matrix:

$$\begin{pmatrix} R_{sRGB} \\ G_{sRGB} \\ B_{sRGB} \end{pmatrix} = \begin{pmatrix} 3.240454 & -1.537139 & -0.498531 \\ -0.826029 & 1.759877 & 0.023179 \\ 0.043396 & -0.198899 & 1.063208 \end{pmatrix} \cdot \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} \quad \text{Eq 7}$$

An XYZ triple may represent a color that is out of gamut for Rec. 709 primaries. Such a color cannot be represented in sRGB code values between zero and unity: An out of gamut color will transform to an RGB triple where one or more components is outside the range zero to one – that is, a component may be less than zero, or may exceed unity. Ordinarily, you will have to clip such excursions to the range zero to one.

Combining the matrix of Equation 7 with the optimized color correction matrix of Equation 6 (for the D₆₅ illuminant, and no IR cut filter) yields this transform of linear-light values:

$$\begin{pmatrix} \hat{R} \\ \hat{G} \\ \hat{B} \end{pmatrix} = \begin{pmatrix} 0.7659 & 0.7052 & -0.4990 \\ -0.3140 & 1.3283 & -0.1367 \\ 0.0609 & -0.4739 & 1.0326 \end{pmatrix} \cdot \begin{pmatrix} R_{TCS230} \\ G_{TCS230} \\ B_{TCS230} \end{pmatrix} \quad \text{Eq 8}$$

Following this transformation, sRGB gamma correction – approximately a 0.45-power function, similar to a square root – is applied to the linear RGB values to obtain $R'G'B'$. See Poynton's book, cited in the margin of page 5, for details.

The optimum 3×3 matrix depends upon the TCS230 responses, which are published by TAOS, and upon the spectral absorbance of optical components such as the lens and any IR filter. In most applications, reflected color is measured; the optimum matrix also depends upon the spectral content of the illuminant, and upon the spectral reflectance of the samples to be estimated. Because these latter factors depend upon the application, the system designer will probably have to compute his or her own optimized matrix. In the example presented above we have used the ColorChecker as a proxy for colors that might be encountered in an application that deals with a wide array of color stimuli. If you intend to use the TCS230 to distinguish among a small handful of different colored items, the linear matrix should be optimized for the spectral reflectances of just those items. If you intend to use an illuminant having spectral properties different from CIE D₆₅, then you should optimize for that illuminant.

The TCS230 evaluation module is equipped with a pair of Nichia white LEDs to illuminate the sample. The evaluation module mounts the TCS230 device behind a lens that incorporates an IR cut filter whose response is shown in Figure 11. Here is the optimum matrix to produce XYZ for the twenty four ColorChecker patches as reported by the evaluation module:

$$\begin{pmatrix} \hat{X} \\ \hat{Y} \\ \hat{Z} \end{pmatrix} = \begin{pmatrix} 0.3895 & 0.4993 & -0.0491 \\ 0.1212 & 0.8890 & -0.1231 \\ 0.0343 & -0.2657 & 0.9438 \end{pmatrix} \cdot \begin{pmatrix} R_{TCS230} \\ G_{TCS230} \\ B_{TCS230} \end{pmatrix} \quad \text{Eq 9}$$

The average magnitude of the error of this transform across the 24 patches of the ColorChecker is about 6 delta-E units.

To produce linear sRGB values from a target illuminated by the white LED and sensed through the IR cut filter, use this transform:

$$\begin{pmatrix} \hat{R} \\ \hat{G} \\ \hat{Z} \end{pmatrix} = \begin{pmatrix} 2.781 & 0.4407 & -0.8524 \\ -0.4221 & 2.377 & -0.1920 \\ 0.1366 & -0.8522 & 2.3492 \end{pmatrix} \bullet \begin{pmatrix} R_{TCS230} \\ G_{TCS230} \\ B_{TCS230} \end{pmatrix} \quad \text{Eq 10}$$

As mentioned on page 14, ordinarily you will have to clip the *RGB* values that result from this calculation to the range zero to one.

TAOS has developed an Excel spreadsheet – or more properly, an Excel workbook – that implements the calculations outlined in this white paper. The spreadsheet is available from TAOS, and is described in the companion document *TAOS TCS230 Spreadsheet description*. ■