Colorimetric Characterization of Digital Camera Sensors

Douglas A. Kerr

Issue 1
October 18, 2015

ABSTRACT AND INTRODUCTION

Digital camera sensors typically have three sets of photodetectors, with differing spectral responses. When a small area of the sensor receives a certain light, the sensors in the three sets (referred to as "channels" in this context) deliver three output values. It would be nice if this set of three values would consistently tell us the color of the light, but for the sensors we commonly encounter, it doesn't

That makes it essentially impossible to, from a set of these values for a pixel of the image, accurately determine what color (under the representation defined for a certain color space) to record for the pixel.

It also makes complicated the matter of describing the response of the sensor (as we might look to some testing laboratory to do for us).

In this article, I discuss these interlocking issues.

An appendix discusses the reports of sensor behavior for various digital cameras published by a well-respected testing laboratory (DxOMark) and discusses some conundrums in them.

1. BACKGROUND

1.1 Terminology and notation

The three groups of photodetectors in a digital camera sensor are often identified as "R", "G", and "B". This is, in a way, reasonable, given that the spectral responses of these three kinds of photodetectors emphasize spectral regions that we can think of, in a comfortably vague way, as "red" green", and "blue".

But I do not use that notation in my technical writings, because these designations can suggest (erroneously) that these three kinds of photodetectors somehow correspond to the three coordinates of some RGB color space (designated R, G, and B) and, as we will see shortly, they do not.

So I (arbitrarily) call the three groups of photodetectors "D", "E", and "F".

Copyright © 2015 Douglas A. Kerr. May be reproduced and/or distributed but only intact, including this notice. Brief excerpts may be reproduced with credit.

It is common to refer to the outputs of the three kinds of photodetectors (often said to be the outputs of the sensor's three "channels") R, G, and B. Of course, I don't like that, in part for the reason I just explained. But further because, in technical wiring, the upper-case letters imply the "nonlinear" form of the coordinates of an RGB-family color space (after application of what is sometimes called "gamma precompensation"). Lower case letters (for example, r, g, and b) are often used to imply the basic, linear form of those coordinates. And when, in this article, I speak of those basic, linear coordinates of a color in, for example, the sRGB color space, I use the symbols r, g, and b.

And accordingly I call the three kinds of outputs from a sensor, which are in linear form, d, e, and f.

1.2 Color, spectrums, and metamerism

Light with a certain spectrum[1] will have a certain color. Sometimes we say, "will appear to a human observer to have a certain color", but in fact, color is defined in terms of human perception, so if a certain instance of light appears to have a certain color, it in fact, by definition, **has** that color.

But in fact, there are an infinity of other spectrums that have that same color. This situation is called *metamerism*, and the (infinite) collection of spectrums that have a certain color are said to be the *metamers* of that color.

1.3 The color filter array (CFA) sensor

In most of the sensors with which we will be concerned, there is not a D photodetector, an E photodetector, and an F photodetector at each pixel position. Rather, photodetectors of the three types are interleaved, there being at each pixel location only one, of one of the three types. Often this interleaving is done with a Bayer pattern, in which in a square cluster of four pixel positions there is a D photodetector, two E photodetectors, and an F photodetector.

An interpolation process called *demosaicing* gives us with d, e, and f value for each pixel location, in each case on of the three having come from a photodetector and the other two being "estimated".

---

[1] The rigorous term is "spectral power distribution function". The electrical engineer would call essentially the same thing the "power spectral density function" (although there are subtleties of difference).

I will not discuss this further in this article, and when I speak of the sensor d, e, and f values from a "small region of the sensor", it will be as if all actually came from directly from photodetectors.

1.4 Colorimetric and non-colorimetric sensors

We often, as one of the ways we "take delivery" of an image from a digital camera, do so in a JPEG file. There, the estimated color of each pixel is recorded in a form derived from an R, G, B representation under the sRGB color space.[2]

It would be wonderful, of course, if the d, e, f outputs were in fact the r, g, b values of the color of the light on the sensor region. But almost never is that so.

Next best would be if the d, e, f outputs told us, unambiguously, the color of the light on the sensor region. Then we could somehow, in a deterministic way, transform a set of d, e, f values to the r, g, b values that also describe that same color.

But sadly, we do not get to enjoy even that nicety for the sensors we normally encounter. Typically, we find that exposing a sensor region to different metamers of the same color will produce different sets of d, e, f values.

So how can we reliably take a set of d, e, f values from our sensor and determine the color of the light of interest, so we can record that color in terms of r, g, b? Well, we can't!

Sensors that behave this way (all the ones we normally encounter in digital cameras) are said to be *non-colorimetric*, which means "does not measure color". (Well, that's pointed enough!)

Of course, it seems to us in our photographic work that camera manufacturers must have found a way to overcome this imperfection. Well, the best they have done is to find out how to minimize its impact on our work. We will look more into that in a later section.

Because the basic syndrome here is that the sensor does not respond consistently (in terms of its outputs, d, e, and f) to different metamers of a given color, the resulting discrepancy is called *metameric error*.

---

[2] In a variant, we may have the JPEG file be based on an alternate color space, Adobe RGB (which may be advantageous as it has a larger color gamut). The does not affect the principles involved here, and I will ignore this alternate possibility for convenience.

2. THE TRANSFORMATION MATRIX

2.1 For a trivial colorimetric sensor

Suppose had a sensor for which we know that the d, e, f values are in fact exactly the r, g, b values for the color of any light onto the sensor. It is the tidiest form of a *colorimetric sensor*. We could consider "transforming" a set of d, e, f values from that sensor (perhaps for a pixel location) to the corresponding r, g, b values this way:

r = d
g = f
b = g

That seems pretty silly for such a trivial operation, but I am actually getting us ready for further, less trivial, work.

Now to get even sillier, we could do that trivial process by matrix multiplication, this way:

$$\begin{bmatrix} r & g & b \end{bmatrix} = \begin{bmatrix} d & e & f \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \qquad (1)$$

That is, we take the triplet d, e, f, treated as a *vector* (a matrix with one row, or one column, in this case one row), multiply it[3] by a 3×3 "identity" matrix (the "trivial" matrix[4]), and get our result as the vector r, g, b.

Now sometimes we show this process in what I call "engineering" form, as it is especially beloved to electrical engineers. That looks like this in the present case:

	r	**g**	**b**
d	1	0	0
e	0	1	0
f	0	0	1

Figure 1.

---

[3] To be precise, "right multiply" it by the 3×3, a subtlety that is important because matrix multiplication is non-commutative; that is, if we reverse the order of the two things being multiplied, we get a different result.

[4] It has the property that when we multiply another matrix by it, the result is the original matrix (just like multiplying by 1 in ordinary algebra).

This considers the matrix to be a linear "network" with three inputs (d, e, and f) and three outputs (r, g, and b). And this layout helps us to remember how the actual calculation of matrix results proceeds. For each output, we take each of the inputs, multiply it by the numerical "coefficient" in its row and in the column of that output, and add all those products (for all the inputs) together to get the output value.

And in turn, this shows how each of the inputs influences the value of each of the outputs. In this very simple case, we see, for example, that input d influences output r (by a factor 1) and does not influence output g (or, we could say, influences it by a factor 0).

2.2 For a less-tidy colorimetric sensor

Now, let's take the case where the sensor outputs d, e, and f do not directly correspond to the r, g, and b values that describe the color of the light on the sensor, but nevertheless consistently describe the colors (the "second nicest" kind of sensor, still colorimetric). In the most desirable case (which I will assume), this means that there is a linear transformation from d, e, f to r, g, b. And this time, to jump right to the punch line, we can do this with a matrix.

This time I won't use specific numerical values, but will use a symbolic representation of the coefficients of that matrix (and I will show it in "engineering" form).

$$
\begin{array}{c|ccc|}
 & \mathbf{r} & \mathbf{g} & \mathbf{b} \\
\hline
\mathbf{d} & T_{dr} & T_{dg} & T_{db} \\
\mathbf{e} & T_{er} & T_{dg} & T_{eb} \\
\mathbf{f} & T_{fr} & T_{fg} & T_{fb} \\
\hline
\end{array}
$$

Figure 2.

We will call this matrix T ("transform") and so I label its nine coefficients with the letter T plus a subscript. Coefficient T_{dr}, for example, is the coefficient that tells by what factor does input d affect output r.

Now, how might we determine, in a laboratory, what the coefficients of this matrix should be? Well, there is only one kind of light whose color's sRGB r, g, b representation is r, 0, 0 (that is, its only non-zero coordinate is “r”): light whose color is that defined for the sRGB “R” primary.[5]

So we can expose our sensor to, in order, “unit values” of the sRGB primaries R, G, and B, for each of them noting what d, e, and f outputs we get from our sensor. We get nine numbers from those tests.

We can present that sensor behavior as a matrix, which (again in “engineering” form) would look like this:

$$
\begin{array}{c|ccc|}
 & \mathbf{d} & \mathbf{e} & \mathbf{f} \\
\hline
\mathbf{r} & T'_{rd} & T'_{re} & T'_{rf} \\
\mathbf{g} & T'_{gd} & T'_{ge} & T'_{gf} \\
\mathbf{b} & T'_{bd} & T'_{be} & T'_{bf} \\
\hline
\end{array}
$$

Figure 3.

We will call that matrix T′ (you will see why in a little bit). The coefficient I have designated T'_{rd}, for example, tells us how the d sensor output is influenced by the “r-ness” of the light on the sensor. We determined that directly in the tests by noting what value of d we get when we apply one unit of light that only has a “r” aspect to its color, namely light whose color is that of the sRGB “R” primary.

Now it turns out that if we take this matrix, T′, and take its inverse (not a trivial process, but there is a straightforward procedure for doing it), that will be the matrix we should use to transform a set of d, e, f values to the corresponding set of r, g, b values. We will call that matrix T. This is consistent with what we did earlier, and will fit best with what we will do later.

Now, if T′ is the inverse of T, T is the inverse of T′. And so it will be legitimate to call the matrix we started with in this case T′.

2.3 But our sensors are non-colorimetric

That is all a lovely, tidy story. But I said earlier that our sensors are not colorimetric. That means that a set of d, e, f values from the sensor does not explicitly imply the color of light onto the sensor

---

[5] Note that no spectrum is defined for an sRGB primary, only its color. That doesn’t hurt us here, since we have stipulated that the sensor we are speaking of at the moment responds to the color of the light upon it, regardless of its spectrum. But that issue rears its head later.

(however we might want to describe that), and thus it cannot explicitly imply a set of r, g, b values that rigorously describe that color.

But suppose we ignore that and use the procedure above, based on measurements of the response of the three sensor channels to the three sRGB primaries, take the matrix made up of those nine results, take its inverse, and use that to transform d, e, f from the sensor into r, g, b.

Well, we really can't make such a test, because:

- No specific spectrum is defined for each of the sRGB primaries, only a color. And there are an infinity of different spectrums that would have that color.

- As this sensor is assumed to be "non-colorimetric", it would in general have different output values for different spectrums having the same color.

So the results of our test would differ according to the actual spectrums our "sRGB primary" test lights happened to have.

But supposed we nevertheless forged ahead (using some handy sRGB primary test lights), made the nine measurements of sensor response, used the results to form matrix T', and inverted that to form matrix T. We then built that into our camera to be used to transform the d, e, and f outputs from the sensor (for example, for some pixel location) into r, g, b values, which we will then record as describing the color of that pixel. How would that work out?

Well, it works out that, in the special cases where the light on a region of the sensor (from our "scene") happened to have the exact spectrum of one of our "sRGB primary" test lights, the resulting r, g, b values would correctly describe the color of that

But otherwise, in general the r, g, b values we got for a pixel will not describe the color of the light on that region of the sensor. That is what it means the sensor is non-colorimetric! We can also say that, for many colors of light, it is subject to *metameric error*.

So what can we do? The best we can do is to transform the d, e, f outputs of our sensor to r, g, b with an arbitrary matrix designed such that, "overall", the resulting error is the least possible.

But what do we mean that "overall, the resulting error is the least possible", and how might we find that "best" matrix?

Well, one set of answers to those questions is found in an ISO standard, ISO 17321. Curiously enough, it does not supposedly deal with the design of camera image chains Rather, it discusses characterization of digital camera sensors. And a major thrust is the development of a “score” for the least degree of overall metameric error we can expect from the sensor involved.

Recall I said that metameric error cannot be overcome; we can just minimize its overall impact under certain circumstances and a certain definition of “overall impact”.

The process defined by ISO 17321 for developing an “overall metameric error score” for a sensor works essentially this way:

- We take a shot with a camera using the sensor of interest of a test chart that includes multiple “colored” patches having well-documented reflective spectrums.

  A typical chart used for this purpose is the X-Rite “ColorChecker” chart. It has 18 “colored” patches, which are used in the ISO 17321 procedure, plus 6 with a “neutral” (white/gray) reflective spectrum, which are not used in the procedure.

  The shot is to be taken with the target illuminated by a certain illuminant[6] (that is, an illuminant with a specific documented spectrum).

- For each patch, we multiply the reflective spectrum of the patch by the spectrum of the illuminant[7], which gives us the spectrum that the light reflected from the patch would have.

- For each patch, we determine, from the spectrum of the light we know would be reflected from the patch, the color of that light (in terms of the coordinates X, Y, Z of the CIE XYZ color space, the description of color used in “scientific” work).[8]

---

[6] This is critical as the phenomenon of metameric error depends on the spectrum of the illuminant under which the objects are photographed. CIE illuminant D55 is the default for this procedure.

[7] This means that, for each wavelength over the range of interest, we multiply the value of the spectrum of the illuminant at that wavelength by the value of the reflective spectrum of the patch at that wavelength, the result being the value, for that wavelength, of the spectrum of the reflected light.

[8] There is a deterministic procedure for doing this.

- We next design a transform matrix to be used (in a "virtual camera", only imagined for purposes of this procedure) to transform the sensor d, e, f outputs into X, Y, Z, which will have the following property: the error between (a) the color represented by the resulting X, Y, Z values for each patch and (b) the "known" color of the light reflected from the patch (determined as described above), averaged over all patches, will be the minimum achievable.[9]

  There is a two-stage deterministic procedure for designing a matrix that will have that property from the d, e, f values for each of the patches and the "known" colors of the light from them.

- That average error (said to be the *average metameric error*) is the basis of the "score" of the metameric error potential of the sensor[10], called the Sensitivity Metameric Index (SMI). The SMI is in percent, and if the least average metameric error that can be attained for the sensor (by using the "best possible" transformation matrix) is zero, then the SMI is 100%.

Does ISO 17321 recommend that, in an actual camera, a transformation matrix designed per this procedure be actually used in a real camera with that sensor? No. ISO 17321 does not have anything to do with designing camera image chains. And remember, the matrix developed in the procedure is not an "output" of the procedure—it is only an "intermediate result".

The only "output" of the process is the SMI. And it is quite reasonable, if not particularly rigorous, to think that the SMI tells us "how near to being colorimetric is the sensor?" If the SMI is 1005, then in fact (so far as can be determined from the specific color patches used in the test), the sensor is (absolutely) colorimetric.

2.4 In actual camera design

So, nevertheless, in the actual design of a camera image chain, does the manufacturer typically go through the procedure in ISO 17321 for the sensor to be used, under some illuminant, and then use the "optimum matrix" developed as part of that procedure to actually transform the sensor outputs (d, e, f) to r, g, b as one step of

---

[9] In fact, both the colors to be compared are transformed from their X, Y, Z representations into the L*a*b* color space, and the difference, in terms of the property "Δe", is considered the "error" for this step of the process.

[10] I say "potential" since, if the matrix used in an actual camera were not the "optimum" one developed in this process, the metameric error performance would be worse than indicated by this "score".

generating an image represented (at one stage of the process) in sRGB form? Well, maybe so, maybe not.

Do they use different "optimum" matrixes, determined per ISO 17321 for different illuminants, based on what they may assume is the illuminant in force from the "white balance" setting of the camera? Dunno.

In fact, more sophisticated cameras may well not use this relatively-simple matrix multiplication process for that transformation at all. Rather, they may use a discrete table that tells, for a large number of discrete d, e, f values, what r, g, b values to transform each of them to (with interpolation used for value combinations that are not in the table).

3. The transform matrix and noise performance

The sensor outputs are afflicted by "noise", that is by a random variation in the output about a value that is the output value the sensor "ought to give".

It is well recognized (and discussed *ad nauseam* in the literature) that if there is substantial asymmetry in the transform matrix (I'll decline here to discuss just what that means), then, overall the impact of this noise will propagate into the r, g, b values more than if the matrix were more nearly symmetrical.

This matter is beyond the scope of this article, but I mention it briefly because it is one issue that a camera designer might be sensitive to in examining a proposed transform matrix.

4. Working backward

4.1 With a colorimetric sensor

If we had a camera with a colorimetric sensor (unlikely to happen in real life), and we have the optimum transformation matrix the ISO 17321 process gave (which in that case would be "flawless" in its work), and took its inverse, what would we get?

We would get a matrix that described the response of the three sensor channels to the three sRGB primaries (a perfectly valid notion for a colorimetric sensor).

4.2 With a non-colorimetric sensor

If we had a camera with non-colorimetric sensor, and we have the optimum transformation matrix the ISO 17321 process gave (which in that case would a “the best we can do” matrix), and took its inverse, what would we get?

Well, precisely, it means only this. If we take some color representation in terms of r, g, b, and multiply it by this inverse matrix, we get the set of d, e, f values that, if it were transformed by our transform matrix proper, would lead to those r, g, b values. No kidding.

But aside from the piece of trivial mathematical circularity, what else is the significance of this inverse matrix? Does it have an understandable physical significance? Well, not really.

Perhaps we think that inverse matrix would be the matrix that would describe “the response of our three sensor channels to the sRGB primaries, R, G, and B”. But again, we must realize that, as discussed above, in the case of a non-colorimetric sensor there is actually no such thing.

Probably the fairest description of the significance of the inverse matrix in this case is, “The coefficients of the inverse matrix tell us generally what is the response of each of the three sensor channels to light from the parts of the spectrum we can think of as “red”, “green”, and “blue”.

Hold that thought when you get to Appendix A.

-#-

Appendix A

The DxOMark sensor report

DxOMark is an operation of the well-respected French company, DxO Labs. DxOMark publishes an extensive series of reports on digital cameras, focusing wholly on image performance (not at all on things such as available shutter speeds, or battery life, or start-up time, or the shape of the handgrip).

One important section of the report is on the performance of the camera's sensor, and one subsection is on its "color response". For reasons that you should by now be alert to, these reports come in two flavors, one pertaining to measurements made under CIE illuminant D50 (a "daylight" illuminant) and one under CIE illuminant A (a "tungsten lamp" illuminant).

Here is the "CIE D50" report for the Canon EOS 500D:

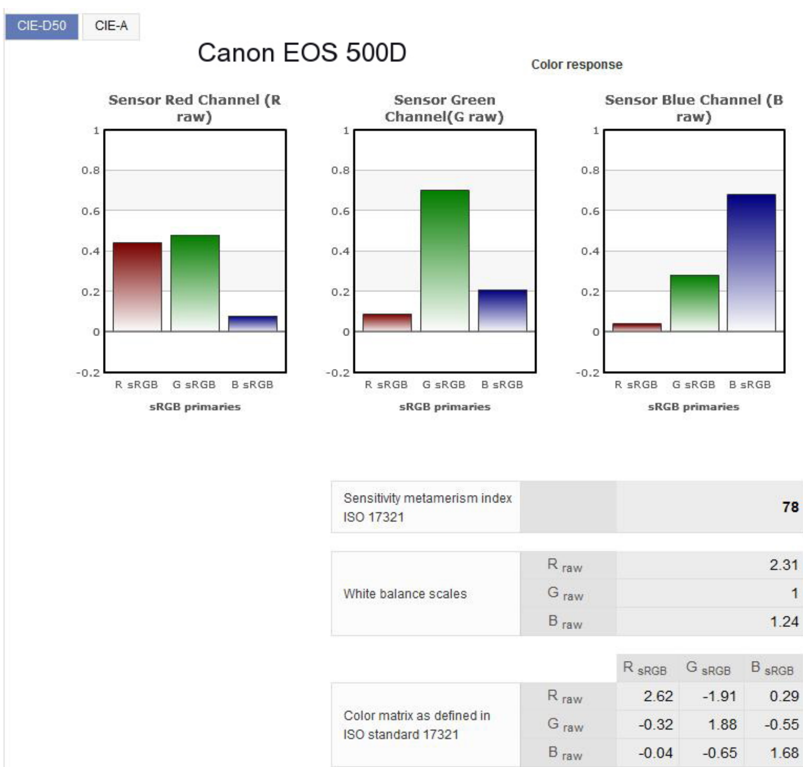

Sensitivity metamerism index ISO 17321		78

White balance scales	R raw	2.31
	G raw	1
	B raw	1.24

		R sRGB	G sRGB	B sRGB
Color matrix as defined in ISO standard 17321	R raw	2.62	-1.91	0.29
	G raw	-0.32	1.88	-0.55
	B raw	-0.04	-0.65	1.68

Figure 4.

Of particular interest to us just now are the leftmost three bar charts and the panel labeled "Color matrix as defined in ISO standard 17321" (well, that latter at least seems familiar).

We might think, from the labeling, that the three bar charts somehow show the response of the three sensor channels (they are identified with the familiar “red”, “green” and “blue”, and “R”, “G”, and “B” notation). The term “raw” tells us that it is the sensor outputs (what I call d, e, and f) for the three channels that are meant.

The labeling of the three bars on each chart as “sRGB primaries” and, individually as “R sRGB”, “G sRGB”, and “G sRGB” suggests that these charts in fact report the results of testing to determine the response of the sensor to the three sRGB primaries.[11]

And the discussion, in some supporting information on the DxOMark Web site sounds, albeit a little vaguely, as if that is what these charts mean.

But of course that can’t be what they mean. As I discussed in the body of this article, for a non-colorimetric sensor (and almost all the sensors of interest are of that type, there is no such thing as “the response of the sensor channels to the sRGB primaries.”

And in fact, DxO has confirmed to me that no, they are not derived from testing of the sensor response to “the sRGB primaries”.

But, I have a suspicion that in some earlier time, that may in fact have been what those bar charts were meant to report. And in their technical explanations, DxOMark seems unable to wholly rid themselves of that notion. (“You can take the *garçon* out of the sRGB primaries, but you can’t take the sRGB primaries out of the *garçon*”.)

Now let’s leave the bar charts for a moment and look at the matrix in the “ISO 17321” panel. Especially given that DxOMark clearly performs the ISO 17321 determination of metameric error on the sensor (since we see the ISO 17321 Sensitivity Metameric Index reported), we might well think that the matrix is in fact the one that is developed in that process.

Well, not exactly. That matrix transforms from d, e, f (my notation—that would be R_{raw}, G_{raw}, and B_{raw} in the notation used in the DxOMark report) to X, Y, and Z, the coordinates of the CIE XYZ color space. From the markings on the matrix in the DxO report (which uses what I call the ”engineering” form), it transforms from d, e, f to “R sRGB”, “G sRGB”, and “B sRGB” (clearly these are what I call r, b, and b).

---

[11] Again we need to recognize that it would be impossible to do that in the context of a non colorimetric sensor, given that no spectrum is defined for an “sRGB primary”.

But there is a rigorous transform from X, Y, Z to sRGB r, g, and b, so it would be quite reasonable that DxOMark had “transformed” the matrix developed under the ISO 17321 procedures into one whose outputs were sRGB r, g, b, and published that.

But I was still vexed by the significance of the three bar charts. Then I had a thought. I took the “ISO 17321” matrix from one of the reports and calculated its inverse (that is, I had Excel calculate its inverse). And whaddya know, the nine coefficients of that inverse matrix corresponded exactly with the heights of the nine bars on the bar charts.[12]

So indeed the bar charts are not derived (except in a very roundabout way) from the sensor characteristics. They are obviously “back formed” from the “optimum” transformation matrix developed as part of the ISO 17321 procedure.

And in fact, DxO has clearly confirmed exactly that to me.

But what is the point, and meaning, of this?

Well, I discussed in section 4. of this article proper that we might expect that the inverse of the ISO 17321 matrix might be broadly similar to a presentation of the “response of the three sensor channels to the sRGB primaries” (recognizing that there is of course actually no such thing).

So maybe the use of the inverse matrix to determine the heights of the bars on the bar charts was a way for DxOMark to satisfy its nostalgia for the postulated earlier notion of using bar charts to report the response of the sensor channels to the sRGB primaries.

But something about that whole story didn’t seem right. And I soon discovered what it was.

If we think in terms of the coefficients of the inverse matrix describing, sort-of, almost, the “response of the sensor channels to the three sRGB primaries”, that would work out if the coefficients of a **column** of the inverse matrix were associated with the bars for one sensor channel, with each coefficient in the column telling (sort-of) the response to the corresponding one of the three sRGB primaries. We can see this from the labeling of the inverse matrix shown in figure 3.

---

[12] In fact the intended numerical values of those bars are on a “tooltip” on each of the bars on the actual report, so I didn’t even need to “read” them on the vertical scale of the chart (I did check for credibility).

But my comparison of the numerical values of the coefficients of the inverse matrix showed that the coefficients in a **row** of the matrix were associated with the bars for one sensor channel, with each coefficient in the row supposedly telling (sort-of) the response of that channel to one of the three sRGB primaries.

So it may well be that, in DxOMark's scheme to keep a little alive the obsolete notion that the bar charts show the "response of the sensor channels to the three sRGB primaries" (as close as there is such a thing for a non-colorimetric sensor), they got the rows and columns of the inverse matrix confused!

Another hint that the "coefficients" have been scrambled as applied to the bars in the bar charts comes from an article on the DxOMark Web site. It presents the spectral responses of the three sensor channels of two cameras, the Canon EOS 500D and the Nikon D5000:

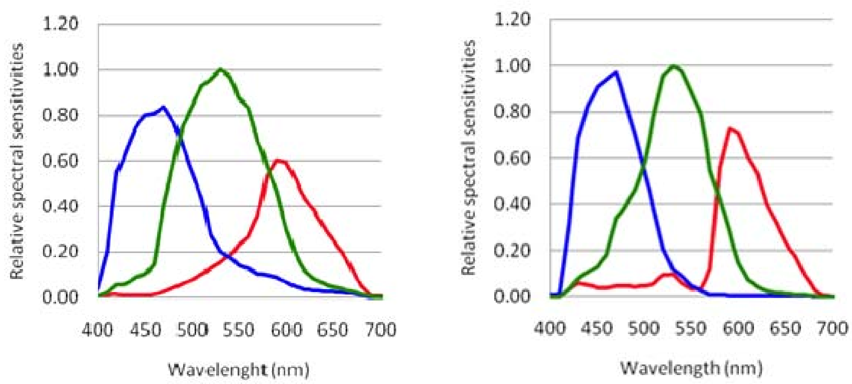

Figure 5.

Attention is directed to the fact that, for the EOS 500D, the "red" sensor channel retains appreciable sensitivity in the region we can broadly think of as "green" (the article specifically notes the spectrum at 550 nm as being "green"). The article points out that this is not advantageous; it represents a lack of "discrimination" in that sensor's "red" channel that will lead to undesirable implications in the processing of the sensor outputs. The intimation is that the Nikon D5000 "does better" in this regard, its sensor's three groups of photodetectors being "more discriminating" as to the portions of the spectrum to which they respond.

Next the article shows us the bar charts from the report for the EOS 500D:

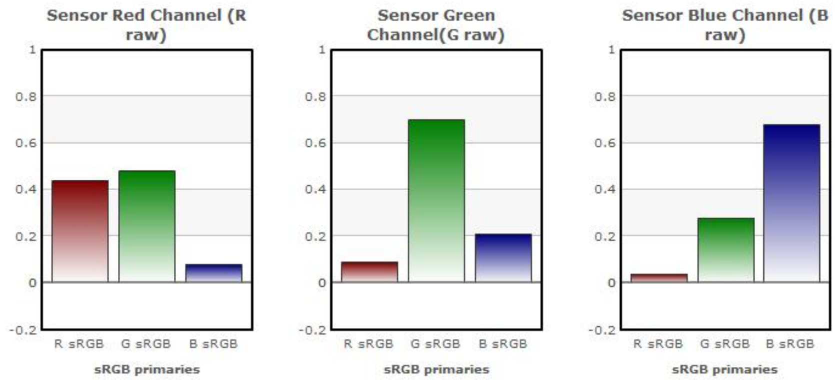

Figure 6.

Excerpted from an image by DxOMark
Used under the doctrine of fair use

Now we know that in fact the heights of the nine bars correspond to the coefficients of the inverse of the "ISO 17321" matrix for this sensor. There is no "precise" physical significance that can be attributed to those coefficients. But they should give us, in the same broad, quantitative sense I referred to earlier, a "hint" of the response of the three sensor channels to "the parts of the spectrum we think of as 'red', 'green', and 'blue' ".

But in fact, the “red” channel chart does not have just a "significant" green bar but rather a very large one—higher than the bar that, broadly, reflects the response of the "red" sensor channel to the "red" part of the spectrum. This does not seem right.

The article also shows the three bar charts for the Nikon D5000:

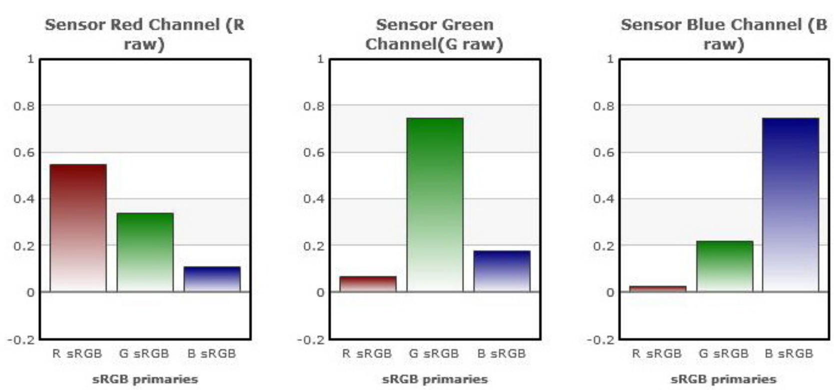

Figure 7.

Excerpted from an image by DxOMark
Used under the doctrine of fair use

Notwithstanding the fact that the article calls attention to the fact that for this sensor the “red” sensor channel has a very small response in the “green” portion of the spectrum, we see in the bar chart a pretty substantial green bar. Again, this does not seem right.

It is the application of the coefficients of the inverse matrix to the heights of the nine bars on the charts that I fear may be incorrect.

Now, if the association of the coefficients of the inverse matrix with the nine bars were done in the way I suspect would be correct, then for the EOS 500D sensor the bar charts would look like this:

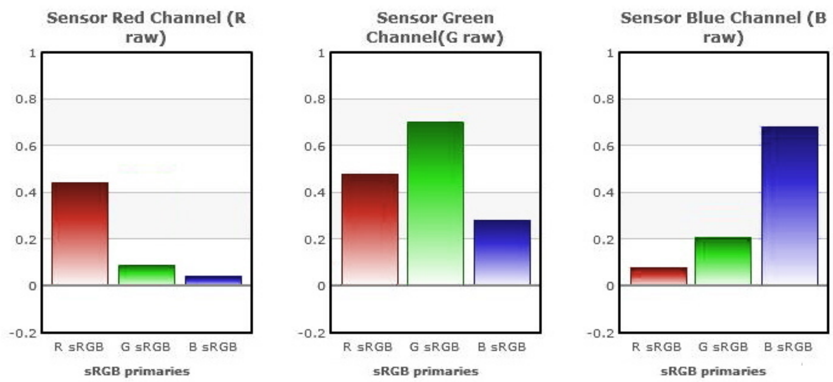

Figure 8.

Adapted from an image by DxOMark
Used under the doctrine of fair use

Now, in the chart for the “red” sensor channel, we see a “significant”, but not large, response to the “green” portion of the spectrum, perhaps what we might guess from the sensor spectral sensitivity chart we saw just above. But there is no simplistic intuitive relationship between the spectral response curves and the coefficients of the inverse matrix, so we can hardly say, “ah, yes, this is clearly correct.”

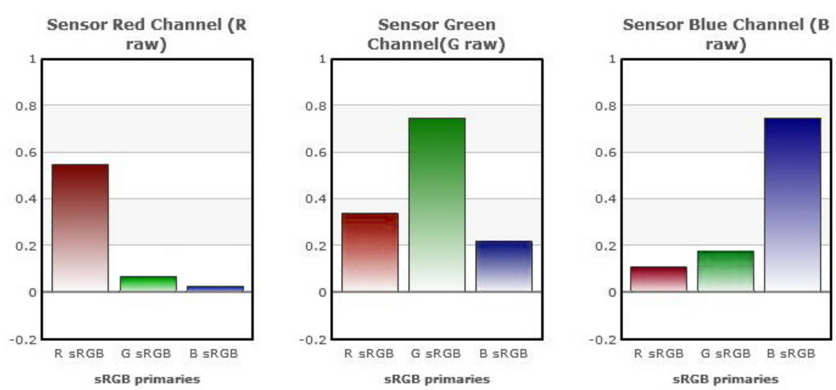

Figure 9.

Adapted from an image by DxOMark
Used under the doctrine of fair use

The figure above shows the bar charts for the Nikon D5000, also revised to match my suspicion as to the correct application of the coefficients of the inverse matrix:

Again this seems to me to be a more credible result considering the spectral sensitivity curves for this sensor. But again, there is no simplistic relationship between the spectral response curves and the coefficients of the inverse matrix, so we can hardly say, “ah, yes, this is clearly correct.”

In any case, this is my current best guess as to the situation:

- Are the bar charts as currently presented in the DxOMark sensor report sensible with respect to how the coefficients of the inverse matrix are assigned to the nine bars? No.

- If that error were corrected, would the presentation be really meaningful? No.

Of course, there are so many ifs, ands, and buts in this story that it is impossible to know for sure what we are dealing with.

I have in any case presented my outlook on this matter to DxOMark. If that leads to any further enlightenment, I will reissue this article to report that.

-#-