Sensor information capacity and spectral sensitivities

F. Cao *[a1], F. Guichard[a], H. Hornung[a], L. Masson[a]
[a]DxO Labs, 3 rue Nationale, 92100 Boulogne Billancourt, FRANCE

ABSTRACT

In this paper, we numerically quantify the information capacity of a sensor, by examining the different factors than can limit this capacity, namely sensor spectral response, noise, and sensor blur (due to fill factor, cross talk and diffraction, for given aperture). In particular, we compare the effectiveness of raw color space for different kinds of sensors. We also define an intrinsic notion of color sensitivity that generalizes some of our previous works. We also attempt to discuss how metamerism can be represented for a sensor.

Keywords: Information capacity, noise characteristics, spectral responses, color sensitivity, metamerism, image quality evaluation.

1. INTRODUCTION

There is a never ending controversy to know whether or not sensor resolution increases image quality. Usually, evaluations on the topic are rather subjective and visually compare photographs taken with different cameras, and conclude that higher resolutions are noisier. This is indeed true if we examine the trend of RAW SNR of recent cameras with smaller pixel pitch. However, a very simple argument makes these studies questionable: from a high resolution camera, it is possible to make a lower resolution camera by suitably down sampling the images. The down sampling basically consists in averaging neighboring pixels and yields an increase of SNR (+3dB for reducing resolution by 50%). With this consideration, it can be seen that the trend is that not only image quality is not getting worse, but is actually getting better. It does not mean that higher resolution is always necessary, it is only use case dependent: printing on a 10x15cm format does not require more than 2Mpix, whereas printing on a 20x30cm does. We aim to bring objective and numerical argument in the discussion. The numerical quantity we bring into the discussion is the information capacity of a sensor, which is the size of the channel that is necessary to transmit an image shot with this particular sensor. Several factors decrease the information capacity of a sensor:

- Spectral responses
- Noise
- Blur, both due to sensor MTF and diffraction blur.

The gray levels on a sensor are encoded on a given number of bits (usually 10 for camera phones and 12 or 14 for DSLR). However, because of the overlap between spectral responses, all these values are physically not attainable on the sensor. As a limiting case, think of three equal spectral responses, boiling down to a monochromatic sensor, but only 10 bits of color depth. Noise changes the effective bit depth, in the sense that two colors that are closer than a noise standard deviation cannot be locally distinguished. Blur changes the effective resolution of the sensor. It is not possible to evaluate the blur of the whole camera since it depends on the lens design. However, two residual causes of blur still remain: blur due to fill factor and cross-talk, measured by the sensor MTF, and diffraction blur that only depends on the lens aperture (through the f-number). In this paper, we discuss the influence of these factors into detail.

[1] fcao@dxo.com

The plan is as follows: in Sect. 2, we focus on the attainable information, describing the set of digital values that can be output by the sensor. The factor into consideration here is the shape of the sensor spectral responses. In Sect. 3, we study the influence of noise. This leads to the notion of color sensitivity. A completely intrinsic notion, only depending on noise curves and spectral responses, is introduced. It correlates very well with the results of some of our previous works, but it is more general. In Sect. 4, we end up with the sensor information capacity that also accounts for blur (related to sensor cross talk, fill factor and lens diffraction). Finally, in Sect. 5, we also discuss how color reproduction accuracy could be represented.

2. ATTAINABLE INFORMATION

2.1 Measuring sensor spectral sensitivities

Sensor spectral responses are measured by using a set of band-pass interferometric filters. Each filter has a bandwidth at half-height of about 10nm. They are backlit by tungsten illuminant. The spectrum of the light out from the filter is measured with a Konica Minolta CS2000 spectrometer. The RAW values of a shot are normalized by this spectrum intensity and result in the spectral response of the sensor.

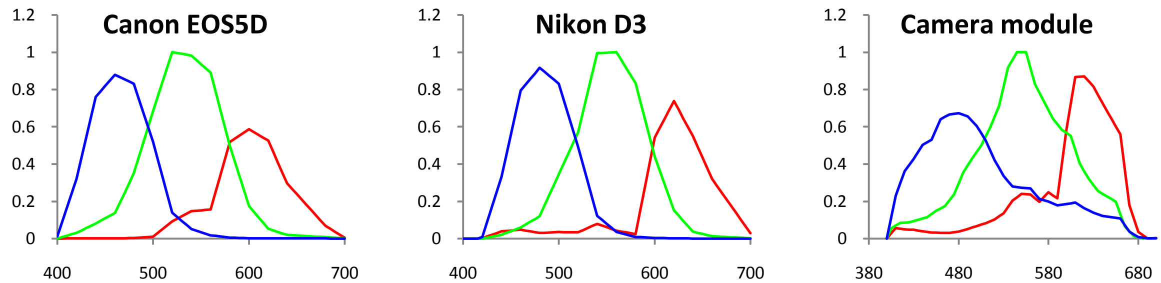

Fig. 1. Spectral responses for two DSLR and a camera phone. The overlap between the different channels is usually much larger for camera phones than for DSLR. Some more subtle differences can be observed between DSLR and are discussed in the text.

Figure 1 shows the spectral responses of two DSLR and a camera module from a cell phone. For this latter, it is worth noticing the very wide spectrum of the blue channel. On the one hand, it increases the sensitivity of the blue channel. On the other hand, in order to obtain images with nice colors, the color matrix will be much more singular since it contains too much green and red. Therefore, we can expect a strong amplification of noise and eventually a possible loss of image quality.

2.2 Effective RAW color space

RAW values are stored on a cube, usually 10 to 14 bits on each channel. However, the whole cube is usually not attainable by the sensor. Indeed, the spectral responses overlap and one channel cannot usually be non zero while the two other ones are zero. If the overlap is very large, a considerable proportion of the cube is actually unused.

This section shows what this proportion is for real sensors. Denote again by r, g, b the spectral responses of the sensor. By definition, these responses directly provide the response to all the monochromatic waves. Since every spectrum can be decomposed on monochromatic waves, the locus of the possible response is the set of positive combination of $r(\lambda), g(\lambda), b(\lambda)$. This set is a cone (denoted by C) with vertex equal to the origin, since changing the exposure simply scales the RAW values by the exposure bias. This cone can be described by its intersection with the plane $R + G + B = 1$. The intersection of the RAW cube with this plane is a triangle whose vertices are the primaries, also called Maxwell's triangle. In this triangle, we can draw the curves of the response of the sensor to monochromatic waves. The possible RGB values belong to the convex hull of the monochromatic response curve. A vertex is attained if and only if two spectral responses are zero for a same frequency.

The ratio between the area of the attainable values and the whole Maxwell triangle describes the efficiency of the sensor in terms of information coding. If there is a huge overlap between the spectral responses, the locus of attainable colors will be very small around a single point, until it collapses into the triangle center (in which case the sensor is actually monochromatic).

It is worth noticing that the behavior between a DSLR and a camera phone sensor is very different, as shown on Fig. 2. Only about 35% of the RAW values are attainable by the camera phone sensor, whereas 80% the typical value for a DSLR. However, for different DSLRs, there are some differences that are particularly noticeable in the medium (green) wavelengths. It can be noticed that the red and blue spectral responses for Nikon D3 are very small at about 560nm where the green channel is maximal. It is therefore possible to saturate the green channel without getting any signal on the other channels. In terms on information coding, this is close to optimal, although the green primary can be approached by only very pure green colors. In the following sections, we try to discuss further on whether it is a good strategy or not.

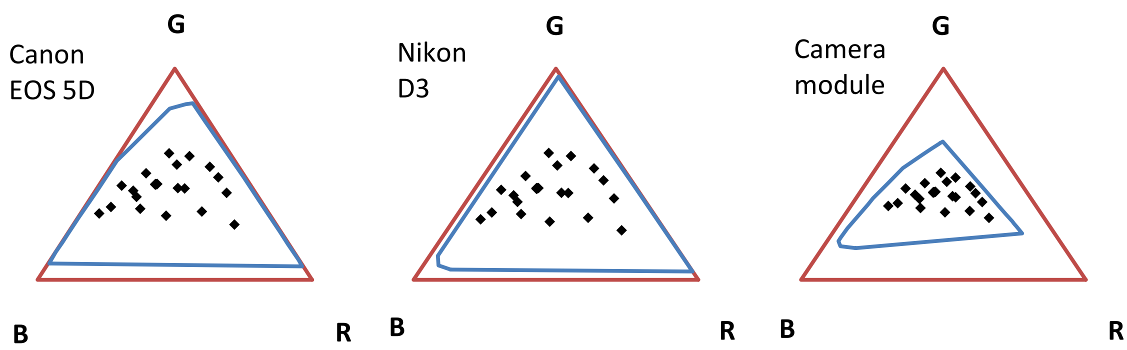

Fig. 2. Effective color space for two DSLR and a camera module (1.75µm, ¼" sensor). The outer triangle is the limit of the possible colors. The vertices are the sensor primaries. The inner curves delimit the set of values that can be attained on the sensor. As an indication, the black dots are the positions of the patches of a Gretag Mac Beth under illuminant D65. They are not used in the calculation of the outer triangle, which is only determined by the response to monochromatic waves.

3. COLOR SENSITIVITY

3.1 RGB Color sensitivity

Color sensitivity[1] was a concept introduced to measure the number of different colors that a sensor can distinguish, up to noise. To be short, two colors in an RGB color space are considered different if their distance is more than one noise standard deviation. The color sensitivity needs the sensor noise characteristics in RAW format. The computation is then as follows. Consider a color rendering composed by a white balance, color matrix. The color matrix is chosen so as to minimize the CIE Lab distance between the Lab values of a chart and the values output by the sensor after this color rendering, as described in the Standard ISO 17321. The noise covariance matrix can be easily computed on the output RGB values from the RAW values. Let $\sigma_1, \sigma_2, \sigma_3$ be the eigen values of the covariance matrix. The color sensitivity in RGB is defined by

$$CS_{RGB} = \int \frac{dr\, dg\, db}{\prod_{i=1}^{3} \max\left(\sigma_i(r,g,b),1\right)}.$$

The integrand can be interpreted as the actual density of discernable colors around the point (r, g, b). Taking the maximum between the eigen value and 1 simply takes the quantization into account.

The maximal color sensitivity in sRGB is 24 bits (each channel is encoded with 8 bits). For a DSLR camera, it is usually about 21 or 22 bits. This is much lower than three times the tonal range (which is usually between 8 and 9 bits for a DSLR camera). We have measured the tonal range and the color sensitivity of most DSLR cameras. These results, as many other ones, are available on www.dxomark.com[2]. As an illustration, Fig. 3 shows DSLR tonal range vs. color sensitivity. Although there is a strong correlation between them, there can be 1 bit difference in color sensitivity for cameras with similar tonal range, showing the importance of spectral responses, via the color matrix.

Like the Sensitivity Metamerism Index (SMI) defined in ISO[3] Standard 17321, color sensitivity makes a clear cut between high end sensors and smaller pitch sensors like those used in camera phones. For a camera phone, color sensitivity is usually about 17 bits. Color sensitivity depends on the illuminant and is more representative of the camera noise than the mere SNR at 18%, since it takes both noise and the sensor spectral response into account. However, there is no need to know the spectral responses since they are only used implicitly. Also, the computed number of colors corresponds to what is seen in the final image. On the other hand, one can always argue that the color matrix depends on the color chart used for the optimization (we usually use a Gretag MacBeth Color Checker), on the illuminant and the color space in which the optimization is performed. Therefore, RGB color sensitivity is not completely intrinsic.

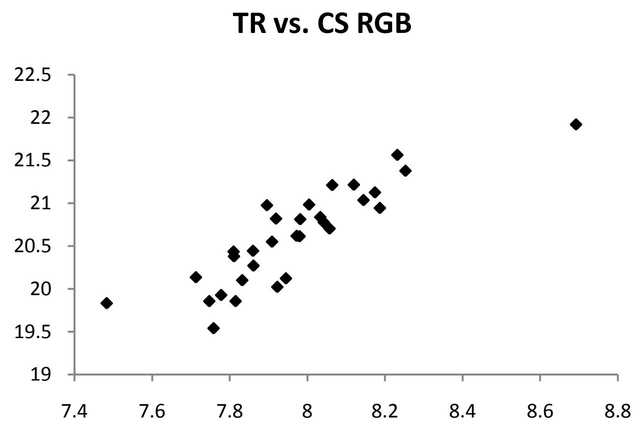

Fig. 3. Tonal range versus color sensitivity for about 30 models of DSLR. There is a strong correlation between tonal range and color sensitivity. However, there are differences of color sensitivity of about 1 bit for sensor with similar tonal range, showing the importance of the color matrix.

3.2 Intrinsic color sensitivity

In this section, we define a completely intrinsic notion of color sensitivity that does not rely on any color chart. However, it requires the sensor spectral responses, and therefore a slightly heavier measurement protocol. From the spectral responses, we can compute the cone C of all possible attainable raw values (see Section. 2) Besides, we also know the RAW noise curves. On all the sensors that we tested, noise is spatially white and its standard deviation only depends on the mean gray value. This might change as the pixel pitch shrinks down, but we make this assumption. The intrinsic (or RAW) sensor color sensitivity is

$$CS_{RAW} = \int_C \frac{dr\, dg\, db}{\max(\sigma_r(r), 1) \cdot \max(\sigma_g(g), 1) \cdot \max(\sigma_b(b), 1)}.$$

It is exactly the same formula as in the RGB case except that the covariance matrix is diagonal. Usually the noise characteristics $\sigma_r, \sigma_g, \sigma_b$ are also equal. Also, the integration domain is not the whole RAW cube but only the attainable part of it. On Fig. 4, we plot the intrinsic color sensitivities for most existing DSLR as a function of manufacturer ISO. The highest one is Nikon D3 (which is a 12Mpix full frame sensor, released in 2007) and the lowest one is Nikon D2X (12MPix APS format, released in 2004). As can be expected color sensitivity drops by about 1.5bits each time ISO sensitivity is doubled (which corresponds to the increase of photonic noise).

There is a very good correlation of the intrinsic color sensitivity with the RGB color sensitivity, as can be seen on the right plot. However, the RAW CS is always about 2 bits higher than the RGB CS. The main reason is that the RGB CS is computed on the sRGB linear color space (i.e. before applying the tone curve). This color space has a small gamut, and because of the large white balance scales on the red and blue channels, about half the values are clipped on these channels (white balance scales more than 2 are not rare), yielding a loss of 2 bits. This shows the dependence to the working color space. Choosing a larger space like Adobe RGB or Prophoto would solve the problem for usual photography. However, the color sensitivity is also meaningful for any type of camera (like multispectral cameras) and should not depend on any particular output gamut.

4. INFORMATION CAPACITY

The information capacity of a sensor[4] is a numerical value that accounts for sensor resolution, pixel MTF and noise characteristic. It is the effective number of bits that can store information on the sensor. It is defined as the product of the effective resolution and the effective bit depth. These two last concepts are defined as follows.

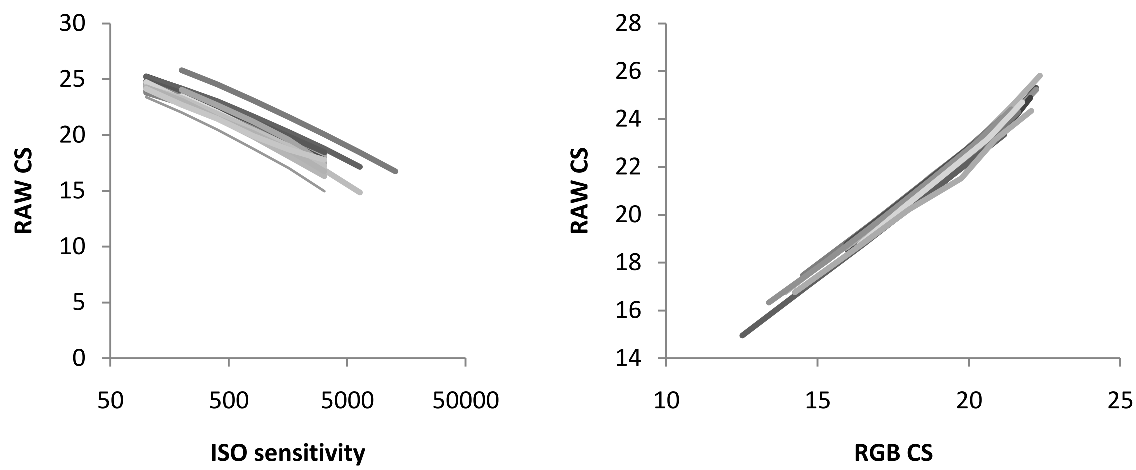

Fig. 4. Intrinsic (RAW) color sensitivity as a function of ISO setting for about 30 DSLR cameras. As can be foreseen, the loss of CS for doubling the ISO is about 1.5bit. Depending on the pixel pitch and the release date, there can bit 5 bits difference between the best and the worst cameras. The RAW CS correlates very well with the RGB CS, but is not bound to any color space nor any particular target.

.

- Effective resolution: the number of pixels of the sensor is only an upper bound of the actual number of small details that can be represented on the sensor. First, there is a loss of contrast in high frequencies due to the sensor fill factor and cross talk, represented by the sensor MTF. This is not negligible since the loss at Nyquist frequency can be more 40%. The loss increases with smaller pitch. Moreover, even a perfect lens has a blur spot due to the diffraction of the aperture. When coming to very small pixel pitch, the diffraction blur spot becomes larger than the pixel, and there is no gain in resolution. Following Shannon's information theory[5], the loss of information in bits can be defined (up to a multiplicative constant) as the integral of the log-MTF. The idea behind the calculation is simple: all frequencies are independent so the loss is the sum of the loss at each frequency. The loss in bits is given by the logarithm (in base 2) of the MTF. For instance, when the MTF equals 0.5, the loss is equal to 1 bit. When summing over the whole frequency spectrum, the total loss is

$$\Delta s = -\int \min(\log_2 MTF(f), bd)\, df$$

  where bd is the bit depth on the sensor and f the bi-dimensional frequencies (in x and y direction). The min in the formula accounts for the fact that it is not possible to lose more bits than what was initially allocated.

- Effective bit depth: gray level is usually encoded on 10 to 14 bits on digital cameras. This is usually larger than the bit depth of the output images (usually 8 bits). A margin is necessary to account for the non linearity of the processing, especially tone curve. However, it is not possible to distinguish 2^{14} gray levels on the sensor, since the data is actually noisy. Buzzi et al.[1] consider that two gray level cannot be distinguished if they are closer than a noise standard deviation. Since the noise standard deviation depends on the mean gray level, the number of distinguishable gray levels on a channel is the tonal range, defined by

$$TR = \log_2 \int \frac{dx}{\sigma(x)}.$$

where the integral is defined on the sensor gray levels.

The information capacity of a sensor with resolution s is

$$C = s \times 2^{\Delta s} \times TR.$$

Let us remark that, up to quantization, information capacity is invariant with respect to convolution. Indeed, denote by K the MTF of the convolution kernel. For each frequency, the noise variance is multiplied by K. Therefore, the tonal range for each frequency increases by $-\log_2 K$, which is exactly the loss of effective resolution.

The problem with this definition is that it does not take the sensor spectral response into account. If a color sensor is deficient and actually monochromatic, it can still have a good information capacity. Therefore, a better definition is obtained by replacing the tonal range by the intrinsic color sensitivity:

$$C = s \times 2^{\Delta s} \times CS.$$

When we study the evolution of information capacity for DSLR cameras, we can indeed observe that sensor resolution is the most discriminating factor. For DSLR pixel pitch, diffraction becomes limiting for small aperture (about f/16). There was also a trend for lower pixel pitch over the years. However technology also improved, and it did more than compensate the loss of sensitivity of each pixel. It was observed that the loss of SNR was about only 70% of what could be expected by a simple scaling of noise characteristics. Therefore, the DSLRs with higher information capacity are full frame sensors with the highest resolution. For camera phones, the situation is slightly different. Indeed, for typical aperture f/2.8, diffraction becomes the main limiting factor for pixel pitch about 1.4µm. Beyond that limit, information capacity does no longer increase. To push this limit further, it is possible to use wider aperture optics. However, for a suitable image quality, it is also necessary to correct some of the optical aberrations yielded by wide aperture lenses, typically spherical aberrations.

5. COLOR ACCURACY

5.1 Exact calculation of fundamental metamers

Metamers[6, 7] usually design reflection objects that are seen the same under a given illuminant and different under another illuminant. More generally speaking, we will say that two spectra are metamers if they are different but cannot be discriminated by the standard observer.

A classical question is to compute the set of colors that are equally seen by the human observer. The calculation is usually done with matrix pseudo-inverse by using discrete values of wavelength. However it is not more difficult with continuous spectra, as we now prove in this section. Denote by $\bar{x}$, $\bar{y}$, $\bar{z}$ the color matching functions (CMF). For any spectrum s, the values output by the standard observer are

$$X(s) = \int \bar{x}(\lambda)s(\lambda)\, d\lambda$$

$$Y(s) = \int \bar{y}(\lambda)s(\lambda)\, d\lambda$$

$$Z(s) = \int \bar{z}(\lambda)s(\lambda)\, d\lambda.$$

Denote by x^*, y^*, z^* the reciprocal basis of the CMF, defined as the unique set of linear combination of the CMF such that

$$X(x^*) = 1, X(y^*) = 0, X(z^*) = 0$$

$$Y(x^*) = 0, Y(y^*) = 1, Y(z^*) = 0$$

$$Z(x^*) = 0, Z(y^*) = 0, Z(z^*) = 1$$

The functions x^*, y^*, z^* are easily determined by solving a 3x3 linear system. The reciprocal color matching functions (RCMF) are not real spectra since they assume non positive values. The fundamental metamer[8] of spectrum s is

$$Ps = X(s)x^* + Y(s)y^* + Z(s)z^*.$$

It is obviously a linear combination of the CMF (since the RCMF are linear combination of the CMF). It is also a metamer of s since it is seen as $(X(s), Y(s), Z(s))$ by the standard observer. For dimensionality reason, it is unique.

It is then very easy to prove that two spectra s_1 and $s_2$yields the same value for the standard observer (i.e. $X(s_1) = X(s_2)$ and similarly for Y and Z) if and only if there is a spectrum t such that

$$s_2 = s_1 + t - Pt.$$

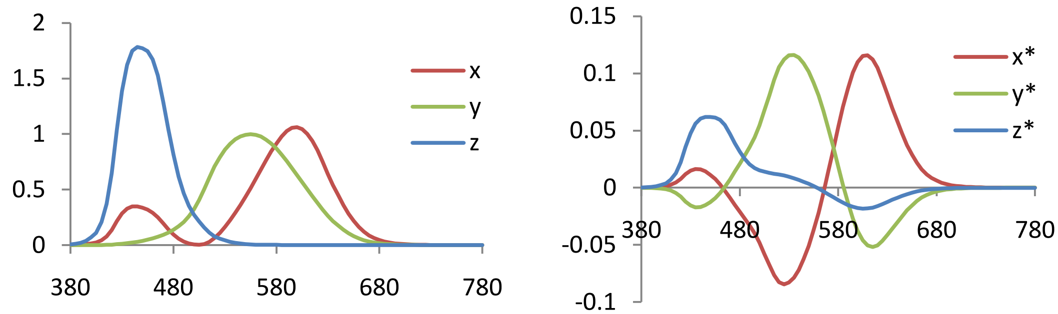

Fig 5. Color matching functions and the reciprocal color matching function (RCMF). They are uniquely determined by some orthogonality conditions. They generate the same space, but RCMF assume non positive values.

Let now r, g, b be the spectral responses of a sensor and $R(s), G(s), B(s)$ the RAW values of the sensor to a spectrum s. The two metamers $s_1$and $s_2$viewed by the sensor differ from $R(t - Pt)$ on the red channel, with similar expressions for the green and blue channels. In particular, it is interesting to decompose t on monochromatic waves. For wavelength λ, the difference between the sensor response to the monochromatic wave at λ and its fundamental metamer is

$$\Delta R_\lambda = r(\lambda) - x(\lambda)R(x^*) - y(\lambda)R(y^*) - z(\lambda)R(z^*)$$

$$\Delta G_\lambda = g(\lambda) - x(\lambda)G(x^*) - y(\lambda)G(y^*) - z(\lambda)G(z^*)$$

$$\Delta B_\lambda = b(\lambda) - x(\lambda)B(x^*) - y(\lambda)B(y^*) - z(\lambda)B(z^*).$$

Every metameric error can be decomposed into a convex combination of these elemental errors.

5.2 Towards an intrinsic measurement of metamerism

In this section we propose a representation of the fact that colors that can be seen equal by the standard observer can yield different raw values on a sensor. The idea is then to examine the error produce by monochromatic waves and their fundamental metamers on the sensor. We add this difference to a chosen illuminant, to simulate the response to a non negative spectrum, although this does not change the result. We then examine in Maxwell triangle the locus of these colors. For a sensor satisfying Luther-Ives condition, this locus is a singleton. Else the size of this set gives an indication on the color reproduction error[9]. It has to be compared with the set of all attainable colors in the sensor.

We made the experiment for the three same sensors as above. The ratio between the locus of the metamers of the illuminant and the possible colors was 1.66% for Canon EOS5D, 3.52% for Nikon D3 and 4.99% for the camera phone sensor. If we look further the shape of the locus of metamers, it can be seen that Nikon D3 has a larger error in the green wavelengths, where the blue and red channels are zero or close.

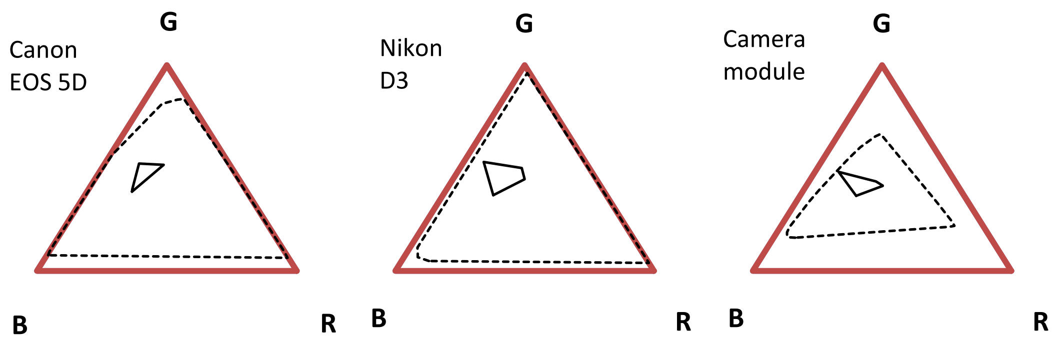

Fig. 6. Values generated by metamers of D65 illuminant on different sensors. The dotted curves are the effective RAW space. The smaller curves are obtained by the spectra obtained by perturbing the illuminant with 5% metamers.

6. CONCLUSIONS AND PERSPECTIVES

In this paper, we discussed different factors that reduce the actual information that can be read from a RAW image. The different contributions yield a number which is the information capacity of a sensor. There are several factors influencing the information capacity of a sensor. We summarize them in the table below.

	Name	**Influencing factors**	**Definition**
Color bit depth	Theoretical bit depth	Memory registry	Memory allocated to store an image
	Effective RAW space	Spectral responses	Subset of memory physically attained by real light signal
	Intrinsic color sensitivity (effective bit depth)	Ditto+noise curves	Number of RAW values that can be distinguished up to noise
	RGB color sensitivity	Ditto+output color space	Number of output RGB values that can be distinguished up to noise
Resolution	Theoretical resolution	Sensor design	Number of photosites
	Effective resolution	Ditto+sensor MTF+ diffraction MTF (lens aperture)	Number of geometrical elements that can be distinguished in the image
Information	**Information capacity**	All	Number of useful bits on the image: product of number of bits per pixel (color sensitivity) and the effective resolution
Color accuracy	***Accurate bit depth***	*Spectral responses*	*Color that can be reconstructed without bias*

Each of these notions can be counted as a number of bits. Effective RAW color space takes into account the overlap between spectral sensitivities and the efficiency of a mere (R, G, B) encoding. Color sensitivity adds the influence of noise, basically by imposing a color quantization of the magnitude of the noise covariance. This gives a number of bits per pixel. Further, information capacity takes the sensor resolution as well as pixel pitch into account by introducing an effective sensor resolution. Smaller pixel pitch is a gain for values larger than typically 1.4µm, but effective resolution is eventually limited by diffraction blur. The product of the effective color depth and the effective resolution is the information capacity.

The next problem is the accuracy of the stored information. Because Luther-Ives conditions are in practice never satisfied, there is an intrinsic bias in color reproduction. We gave some hints how to quantify this inaccuracy, but there is still some work to formulate the color accuracy in terms of information capacity. The ultimate goal of this work is to find the best way to optimize sensor spectral responses[10]. To this purpose, the best trade-off between color accuracy and sensor sensitivity has to be found. This requires being able to compare a stochastic error (noise) and a systematic bias (due to Luther-Ives conditions). Moreover, the information capacity is usually much smaller than the theoretical capacity. There are two possible ways to optimize a sensor: tune every component to increase information capacity. Another possibility is to directly take all the losses into account and only store the useful information. For instance, storing the red and blue channels on the same number of bits is often useless, since the higher bits are usually lost by applying white balance scales. All in all, co-designing the sensor and the image processing would be beneficial for the camera.

REFERENCES

[1] Buzzi, J., Guichard F., Hornung H., "From spectral sensitivities to noise characteristics," Electronic Imaging, 2007.

[2] www.dxomark.com, "A free resource dedicated to RAW-based camera image quality."

[3] ISO 17321, "Graphic technology and photography — Colour characterisation of digital still cameras (DSCs) — Part 1: Stimuli, metrology and test procedures."

[4] Tisse, C.L., Guichard, F., Cao, F., "Does resolution really increase image quality," Electronic Imaging, 2008.

[5] Shannon, C.E., "A mathematical theory of communication," The Bell System Technical Journal 27, 623-656 (1948).

[6] Sharma, G., "Digital Color Imaging Handbook," CRC Press, 2003.

[7] Olerari, C., "Metamerism," Il nuovo Cimento, 9(8): 979-985, 1987.

[8] Cohen, J.B., Kappauf, W.E., "Metameric color stimuli, fundamental metamers, and Wyszecki's metameric blacks," J. Psychol. 95(4): 537-564. 1982

[9] Horn, B., "Exact Color Reproduction," CVGIP, 26: 135-167, 1984.

[10] Quan, S., "Evaluation and Optimal Design of spectral sensitivities for digital imaging," PhD Thesis, RIT, 2002