Sensor spectral sensitivities, noise measurements and color sensitivity

Frédéric Cao, Frédéric Guichard and Hervé Hornung

[a]DxO Labs, 3 rue Nationale, 92100 Boulogne-Billancourt FRANCE

ABSTRACT

This article proposes new measurements for evaluating the image quality of a camera, particularly on the reproduction of colors. The concept of gamut is usually a topic of interest, but it is much more adapted to output devices than to capture devices (sensors). Moreover, it does not take other important characteristics of the camera into account, such as noise. On the contrary, color sensitivity is a global measurement relating the raw noise with the spectral sensitivities of the sensor. It provides an easy ranking of cameras. To have an in depth analysis of noise vs. color rendering, a concept of Gamut SNR is introduced, describing the set of colors achievable for a given SNR (Signal to Noise Ratio). This representation provides a convenient visualization of what part of the gamut is most affected by noise and can be useful for camera tuning as well.

Keywords: Color sensitivity, spectral response, signal to noise ratio, gamut, image quality evaluation

1. INTRODUCTION

Colorfulness is a major attributes of image quality. Indeed, it affects all frequencies, and major color failure can be seen at a very first glance even on a small thumbnail. Therefore, it is of utmost importance to evaluate or predict whether a camera is able to have a good color rendering. The quality of a color rendering is by nature very subjective and relies heavily on personal taste, past experience or even cultural preference. Defining precisely what a good color rendering is cannot reach a general consensus, and is definitely out of the scope of this paper. This paper only deals with objective characterization of color rendering by digital cameras, for which the emphasis is usually put on two major factors:

- how rich is the set of colors that a camera can reproduce?
- How accurate are the colors?

The notion of gamut can be introduced to answer the first point, and will be detailed hereafter. The second point needs to be explained. In general, color accuracy is viewed as colorimetric accuracy, regardless of all the imperfections due to the electronics. However, this (in)accuracy can be dominated by another source of errors due to noise. This is particularly true for low-end cameras, as cameraphones, which have a very small pixel pitch (typically 2.2μm or 1.75μm). Moreover, these devices are often used in low light conditions, like bars or night clubs, with typical illumination of 5 or 10lux.

Before going on with more details, it is necessary to start with a description of color rendering on a typical digital camera. Like human vision, a digital camera usually has three different types of photosites, characterized by their spectral sensitivities, representing the response of the camera to each wavelength. Since they are centered on respectively large, medium and small wavelengths, they are generically called r, g and b for red, green, blue. To make things simpler, we integrate all the different components of the camera into these spectral responses: these usually include the transmittance of the lenses, the infrared filter, the spectral response of the color filter array, and the response of the silicon. At a given gain, the expected response of the red channel to an illuminant I reflected by an object with reflectance E is

$$R = \beta \int I(\lambda)E(\lambda)r(\lambda)\,d\lambda + \delta_r, \tag{1}$$

Further author information: (Send correspondence to Frédéric Cao)
Frédéric Cao: E-mail: fcao@dxo.com, Telephone: +33 (0)155205587

where β is a multiplicative factor (overall gain) and δ_r is an offset due to the electronics. The same holds for the green and blue channels (G and B values). The offset can be compensated, and with no loss of generality, it is set to 0. The factor β depends on the exposure and the different gains. Not only does it multiply the sensor values by a scalar, but it also multiplies the noise by the same factor.

Now, the raw RGB values change from one sensor to another since each sensor has its own spectral responses, which are usually quite different from the primary colors of the device used to display the images. The minimal set of color transformations used to adapt the sensor to the display are

- white balance to compensate for the illuminant.
- chromatic adaptation fitting the sensor color space to the output device color space (or a normalized, device independent color space, the final conversion being done by the output device)
- tonal curve, historically used to compensate for the nonlinearity of CRT screens, and also used to tune the contrast.

The white balance is usually determined by two gain factors applied on the red and blue channels, the green channel being taken as a reference. After white balancing, an object with a neutral reflectance should appear essentially with $R = G = B$, although it can be deliberately set to be slightly different from this. The chromatic adaptation is usually modeled as a 3×3 matrix letting the vector $(1, 1, 1)^t$ be invariant. More complex models (using 3D lookup tables) are possible, but we will always use the matrix model in the following, and refer to it as the color matrix. We will ignore the tonal curve, since all the measurements described in the following are performed before application of the contrast change, or actually require to inverse the tonal curve. The white balance scales are crucial since they can lead to a global and unnatural color shift in a picture. To an extent, the color matrix determines how accurate, vivid or dull the image appears. The role of the color matrix is also to map the sensor color space to another color space depending only on the display device. Therefore, the comparison of sensors is more adequate in this color space, which is supposed to be the same for all the sensors. Changing the output device or the illuminant also requires a different color matrix. The outline of this article is as follows. In Sect. 2, we will develop the concept of gamut of an input device (already studied in several previous works) and compare different types of cameras. As a result, we will see that even though some differences can be observed between low and high-end cameras, results are mostly conditioned by the choice of the calibration of color rendering. Moreover, the concept of gamut also neglects the noise introduced by the camera to obtain a given color rendering. In Sect. 3, we will introduce the concept of color sensitivity, and explain why this measurement is more discriminating for the quality of a camera. One advantage of color sensitivity is that it leads to one single number, and provides a direct comparison of cameras. However, a more local analysis to determine which colors are most affected by noise can be useful. Section 4 introduces the concept of Gamut SNR, which is the set of colors achievable for a given SNR value. It can be very useful, especially for camera ISP (Image and Signal Processing) tuning, since color rendering is mostly a trade-off between the vividness of colors and noise, and is a key for final image quality. We will display the Gamut SNR of several types of cameras before concluding.

2. INPUT DEVICE GAMUT

2.1 Definition

The concept of gamut has been primarily introduced for characterizing output devices. It is defined as the set of visible colors that the device can render. Although this set can be huge, this is actually a simple problem. Indeed, the colors output by the device are combination of a small set of primary colors. Since the human eye has three types of cones, using three primary colors is usually enough to obtain suitable colors, although the gamut can be sensibly smaller than the set of visible colors, depending on the primaries. It can be enlarged by choosing different and/or more primaries.

The gamut of an input device (such as a sensor) is defined as the set of colors that the device can distinguish. This is much more difficult to determine. Indeed, it would require to measure (or simulate) the response of the sensor to all possible spectra, which form an infinitely dimensional vector space. Even by sampling the wavelengths with a finite accuracy (for instance 10nm between 380nm and 800nm), this still remains an intractable computational challenge. However, practical solutions have been proposed,[1–4] though inevitably approximate. To sum up, the different methods consist in choosing a finite set of spectra that can be representative of all possible colors. This set is crucial since it determines how the chromatic adaptation of the sensor is performed (its color matrix). Different possibilities have been proposed as optimal colors, such as the Munsell book of colors samples, or the Gretag Macbeth color checker, although each method has its own limit.[5–7]

In this paper, the purpose is to compare the performances of different sensors. Although a different calibration of a sensor yields a different gamut, two protocols at least can be applied.

- apply the same chromatic adaptation method (same set of color samples, same metric, same illuminant)
- use the color rendering used by the camera manufacturer.

The first method is more objective, although it does not reflect the colors actually output from the camera. Conversely, the second method is subjective, since it reveals aesthetic choices of the camera manufacturer.

2.2 Experimental measurement

The following protocol is used to compute the gamut of sensors.

- Inputs:
  1. Sensor spectral response.
  2. Colorchecker reflection spectra.
- Algorithm:
  1. Compute the raw values of the Color Checker from the spectral response of the sensor and the reflection spectra (see (1)).
  2. For a color matrix A mapping the sensor color space on CIE XYZ, compute the corresponding CIE Lab values, and the mean related error ΔE on the patches of the Color Checker.
  3. Find the matrix A minimizing the mean ΔE error.
  4. For this optimal matrix, draw the (x, y) values corresponding to the response of monochromatic waves.

Some measurements were performed on selected digital cameras: 2 DSLRs, 2 cameraphones. The calibration matrix is strongly influenced by the choice of the target. The output gamuts are all larger than the sRGB Gamut. There was no guarantee for that, since the patches of the ColorChecker are not particularly saturated. The gamuts of the DSLRs are usually larger than the gamut of the cameraphones but not that much, which shows that the measure is not very discriminative. Also, the mean ΔE on the DSLR is much smaller than on cameraphones. This is related to the difference of metamerism of the sensor, as defined by the ISO Norm 17321.[8] As such, the measurement of the gamut and the ΔE error reveals sensor metamerism, but it remains indirect. However, a more direct measurement would be to determine the set of responses of a sensor that can be seen as a single color by the eye.

2.3 Limitations and conclusions

In conclusion, the concept of the gamut of a sensor is not highly discriminative per se, as far as image quality evaluation is concerned. Indeed, it is extremely dependent on the set of colors used to match the spectral response of the sensor on the color matching functions. By using a simple linear model white balance+color matrix, it is observed that there is a tradeoff between the gamut, which is a boundary problem, and the accuracy, which reflects the colors deep inside the gamut. Now, a camera reproducing accurate RGB values is usually perceived as a bad camera, since people usually prefer pictures with saturated colors. Hence, a calibration targeting color accuracy is not characteristic of the final rendering of a camera. Moreover, it is always possible to use a more

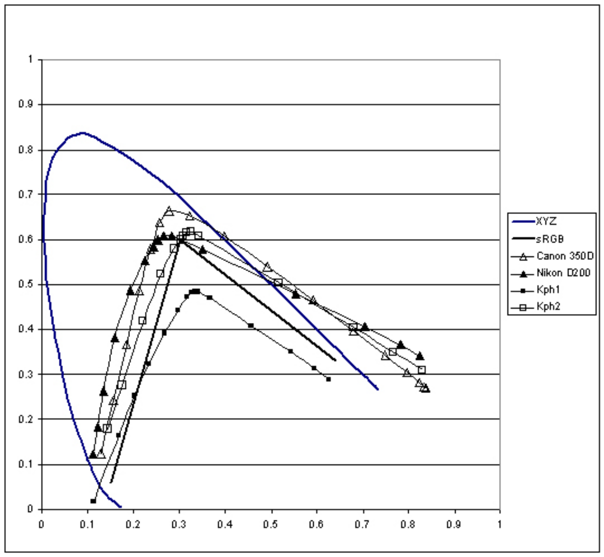

Figure 1. Comparison of the gamut of sensors. Two DSLR and two cameraphones are tested. Since the calibration is performed on the patches of the Gretag Mac Beth in sRGB, the gamuts do not cover much more than sRGB. The gamut of the DSLR is larger, but can still be comparable with cameraphones.

complex transform than a simple color matrix, like a 3D look-up table. The number of degrees of freedom is then huge, and it can be possible to extend the boundaries of the gamut of a camera without sacrificing the accuracy of the inner values. However, there are two problems that a 3D lookup table cannot solve. The first problem is metamerism: if the sensor outputs the same raw values for spectra that are discriminated by the color matching functions, the information is definitively lost. The second problem is the effect on noise: stretching the color space of the sensor to fit a target color space yields an amplification of noise. This is completely ignored by the concept of gamut, and is the main point of the rest of this article.

3. COLOR SENSITIVITY

3.1 Definition

The precise analysis of the colorimetric properties of a sensor is interesting as an index of the theoretical performances of color rendering: accuracy, richness of colors, metamerism problems. This is partially covered by the gamut of the sensor and was discussed in the previous section. However, it is not really representative of the quality of the image that a camera outputs. Indeed, when dealing with real cameras, noise is a crucial factor, especially in low lights conditions, which tends to be a very wide use case for camera as cameraphones. Colorimetric analysis assumes that cameras have an infinite signal to noise ratio (SNR), or that an infinitely wide color patch is observed. Of course, this is unrealistic. Applying a chromatic adaptation matrix (or any look up table) not only transforms the colors of the sensor but it also transforms its noise. For instance, it is clear that the lack of sensitivity in a given channel can be compensated by a gain (which can be analog or digital). Amplifying the signal unfortunately amplifies the noise as well. Chromatic adaptation can also amplify the noise, particularly when the spectral responses of the sensor show a large overlap. Intuitively, the spectral responses of the sensor have to be stretched more to fit the color matching functions. Technically speaking, the color matrix has large singular values.

Hence, another notion of the quality of color rendering has to be introduced, and needs to take the noise of the sensor into account. This is the purpose of color sensitivity, introduced by Buzzi et. al.[9] It is defined as the number of colors that a sensor can distinguish, up to noise. Consider for instance a sensor encoding the gray levels on 10bits on each color channel (which is typical for cameraphones and low-end DSCs). In theory, the sensor can output 2^{30} different values. However, these values are noisy. Noise can be modeled as an additive Gaussian noise. We consider that two values closer than one noise standard deviation cannot be distinguished.

In other words, the actual density of gray levels is the inverse of the standard deviation. In three dimensions, a Gaussian noise is determined by a covariance matrix. The standard deviation is replaced by a confusion ellipsoid. Therefore, there is a limiting color resolution, which we take equal to

$$\frac{1}{\prod_{i=1}^{3} \max(\sigma_i(r,g,b),1)}, \tag{2}$$

where the σ_i are the square roots of the eigen values of the covariance matrix at the point (r, g, b). The denominator is basically the volume of the confusion ellipsoid at the point (r, g, b) bounded by below by the quantization step. When summing this quantity other the whole set of possible values, we obtain the color sensitivity defined by

$$CS = \int \frac{dr\, dg\, db}{\prod_{i=1}^{3} \max(\sigma_i(r,g,b),1)}, \tag{3}$$

the domain of integration being the output color space. The noise covariance matrix can be obtained from measurement on the raw signal, and then transformed by white balancing, color matrix and tonal curve. Taking the log2 of the color sensitivity expresses it as the number of bits encoding the colors on the sensor.

Note that evaluating the color sensitivity does not require a sensor spectral responses measurement. It can be deduced from the noise characteristics and the color rendering only. However, it can be simulated for a sensor whose spectral responses are given, since the raw signal can be simulated as well. Color sensitivity is also much more relevant than the mere raw SNR. Indeed, this latter can be increased by enlarging the spectral responses of the sensor. However, a correct color rendering can only be obtained by substantially degrading the noise by an extreme color matrix.

3.2 Good SNR/bad color sensitivity: a text book case

As an example, let us consider a sensor with a given spectral response and color sensitivity. Let us denote by (R, G, B) the raw values of the sensor. Assume also that the covariance matrix Σ is diagonal, all diagonal terms being equal to σ^2. This sensor has a color matrix, denoted by M. Let us now assume that the spectral responses are extended into a fictive sensor. Let us denote by (r, g, b) the raw values of this sensor, and assume that they are obtained from (R, G, B) by the following relation

$$\begin{pmatrix} r \\ g \\ b \end{pmatrix} = A \begin{pmatrix} R \\ G \\ B \end{pmatrix},$$

where A is the 3×3 matrix

$$A = \begin{pmatrix} 1 & 0.5 & 0 \\ 0.25 & 1 & 0.25 \\ 0 & 0.5 & 1 \end{pmatrix}.$$

Each photosite is 50% more sensitive than on the original sensor. In order to have the same sensor sensitivity, the gain needs to be 66% of the original value. If we assume that the noise is mainly photonic, the raw noise variance has been multiplied by 0.66, which is an SNR increase of 1.76dB. However, the color matrix has to be multiplied by

$$B = \left(\frac{2}{3}A\right)^{-1} = \begin{pmatrix} 1.75 & -1. & 0.25 \\ -0.5 & 2. & -0.5 \\ 0.25 & -1. & 1.75 \end{pmatrix}$$

to obtain the same colors. The new noise covariance matrix is then $MB\Sigma B^t M^t$. Basically, the color resolution has been decreased by a factor $\det(B\Sigma B^t)^{1/2} = 2.44$, which is equivalent to a loss of 1.29 bits. (Here, we neglected the quantization effect, which makes the degradation even worse.) Therefore, even though the sensor has a much better SNR in raw, its color sensitivity has decreased.

3.3 Experimental measurements

The algorithm to compute the color sensitivity of a sensor is the following:

- Inputs:
  1. sensor raw values of the Colorchecker
  2. raw noise curves of the sensor
  3. target values of the Colorchecker in sRGB linear color space.
- Algorithm
  1. Find the color matrix minimizing the mean ΔE error in CIE Lab color space between the target values and the observed values.
  2. For each point in linear sRGB (no gamma curve applied), compute the noise covariance matrix by using the raw noise and the color matrix.
  3. compute the color sensitivity by integration, as given in (3).

The following graphs represent some results of the color sensitivity of 8 cameraphone sensors and 21 DSLR cameras (from old models to the most high-end recent ones) under illuminant D65. In order to show that raw noise does not always mean low noise after processing, the SNR is displayed versus the color sensitivity. The measurements are performed at real ISO 100 (which can be slightly different from the ISO announced by the manufacturer). The SNR is measured on the green channel at 18% of the dynamic (which is a usual target exposure). There is a correlation between SNR and color sensitivity: DSLRs are always better than cameraphones; but cameras with the best SNR do not necessarily have the best color sensitivity. However, at equivalent SNR, color sensitivities may differ by up to one bit. The way ISP manages low exposure/high sensitivity is also interesting. To this end, let us compare noise and color sensitivity measurements at ISO 100 and ISO 1600. If the noise were only photonic, the SNR should decrease by 12dB and the color sensitivity should decrease by 6bits. However, different cameras can have different behavior when the gain (sensitivity) increases, especially in shadows. The color sensitivity takes the global noise behavior into account. Note also that color sensitivity is very dependent on the illuminant. Indeed, in contrast to raw noise, white balance scales and different color matrices have to be taken into account, and eventually give more relevant measures. It is

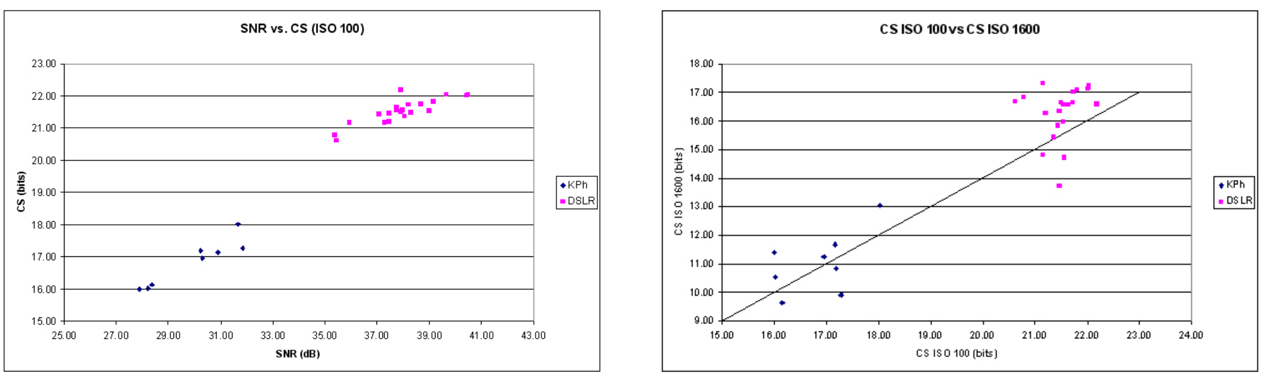

Figure 2. Left: raw SNR vs color sensitivity for different cameras. DSLRs are clearly better than cameraphones (KPh). More interestingly, there are some some ranking inversion between cameraphones when the SNR or the color sensitivity is considered. This latter is more meaningful. Right: color sensitivity at ISO 100 vs. color sensitivity at ISO 1600. If the noise were purely photonic, the color sensitivity should drop by 6 bits (represented by the black straight line). Some sensors beat this limit. This diagram is representative of the performance of the camera in low light.

possible to argue that the raw SNR is not relevant since it does not provide an evaluation of the performance on the final RGB image (after raw conversion). Indeed, any raw conversion includes a denoising algorithm. This is true, but a measurement of noise in a supposedly uniform area is not a perfect measurement either. Indeed, it is very well known that most ISPs smoothen the uniform areas to increase the SNR. However, this creates a large and colored grain in the pictures and is also degrades areas with thin textures. Still, color sensitivity can be used to compared different sensors when using the same ISP.

4. GAMUT SNR

4.1 Definition

The color sensitivity is a good global index that can be provided with an executive summary, since it directly allows direct sensors performances comparison. However, it may be useful to have more local information, and exhibit what part of the gamut is the most penalized. More precisely, it is usually considered that a SNR equal to 10 (that is 20dB) is the minimal value to obtain a correct image, and that an image is good for SNR equal to 40 (i.e. 32dB). Because of the different white balance gain and the color matrix, different parts of the output gamut exhibit quite different SNRs. We choose to represent these values in the CIE Lab color space, for different values of L. It is possible to determine the confusion ellipsoid for each (L, a, b) triplet. The axes of this ellipsoid are given by the noise covariance matrix $\Sigma(L, a, b)$, which is easily calculated from the RAW noise, the white balance scales, the color matrix and the Jacobian of the transformation from XYZ to Lab. An interesting parallel can be drawn with the Mac Adam ellipses: for a given color, it is the set of colors it cannot be distinguished from. The CIE Lab was designed such that these ellipses should be circles with radius equal to 1. Here, we suggest that in addition to the limiting resolution of perception, noise also makes colors indistinguishable. We define the SNR at value (L, a, b) by

$$SNR(L, a, b) = \frac{\sqrt{L^2 + a^2 + b^2}}{\sqrt{\text{trace } \Sigma(L, a, b)}}. \tag{4}$$

For a given threshold τ, the Gamut SNR-τ is the set of values (L, a, b) for which $SNR(L, a, b) \geq \tau$. For a sake a clarity, the Gamut SNR is represented in the ab-plane for different values of L.

The Gamut SNR extends some industry standards whose purpose is to determine the flat field illumination which is necessary to obtain SNR=10 on the luminance (obtained as a linear combination of R, G, B after white balance and color matrix).

For a given value of L, different sensors can be compared. Moreover, the value of the SNR in the ab-plane for a given value of L shows which colors are the most noisy. The values around $a = b = 0$ are usually the most noisy, which is also perceptually relevant, since we are very sensitive to local hue shift in areas that should be neutral.

4.2 Experimental measurements

The method to compute the Lab SNR is as follows

- Input:
  1. raw values of the Gretag MacBeth Color Checker
  2. raw noise curves of the sensor
  3. CIE Lab values of the patches of the Color Checker for the used illuminant
- Algorithm
  1. Determine the color matrix best fitting the sensor raw color space and CIE XYZ for the patches of the Color Checker. The fitting error is the ΔE in CIE Lab.
  2. For each (L, a, b) value corresponding to a (X, Y, Z) value in the visible spectrum, compute the noise covariance matrix.
  3. Compute the SNR by using (4).

Measurements for a DSLR (Canon EOS 400D) are presented on Fig. 3 for luminance values 30, 50, 70 and illuminant D65. Every measurements were performed with gain or real ISO sensitivity 100. (We distinguish the manufacturer ISO corresponding to the camera setting and the ISO sensitivity as defined in the norm ISO 13232[10]). As predicted the SNR increases as the luminance increases. Moreover, neutral colors (in the vicinity of $a = 0$, $b = 0$) are a local minimum of the SNR at a given luminance. This is perceptually consistent: it is well known that saturating an image amplifies its noise, and that it is very conspicuous in neutral areas. Colors in the yellow tones (a close to 0 and positive b) also have a bad SNR because they correspond to a value $Z \simeq 0$ which penalizes the Lab noise.

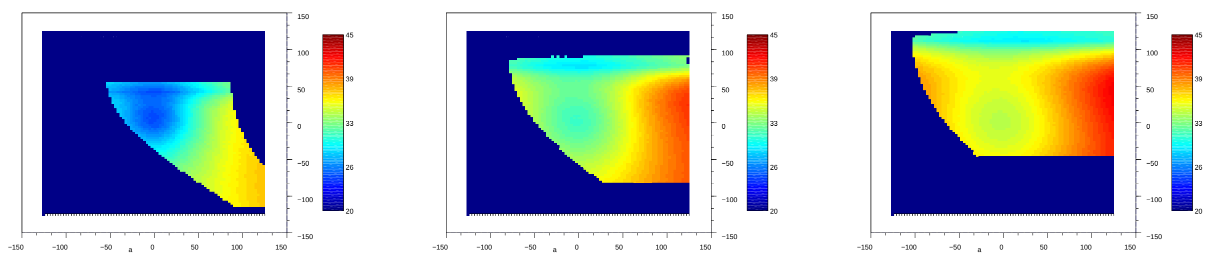

Figure 3. Gamut SNR of the Canon EOS 400D at $L = 30, 50, 70$. The referent illuminant is D65.

On Fig. 4, three DSLR are compared for luminance L=50 (the Canon EOS 400D, Nikon D80 and Pentax K10D). Pentax K10D is clearly the best one. Canon EOS 400D and Nikon D80 have very similar results, and experimental measurements shows that they indeed have the same color sensitivity as well.

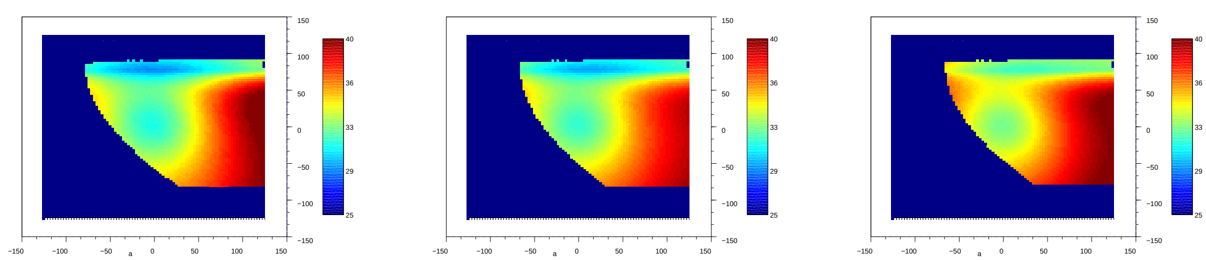

Figure 4. Gamut SNR at $L = 50$. The referent illuminant is D65. From left to right: Canon EOS 400D, Nikon D80, Pentax K10D.

On Fig.5, three camera modules are compared. The two first ones are from the same manufacturer (denoted by M1). The first sensor has a $2.2\mu m$ pixel pitch, the second one $1.75\mu m$. However, the second one is better in terms of color noise, showing that the mannufacturer manages (in this case) to maintain the quality, even though the pixel size goes down. However, the third sensor (a $2.2\mu m$ pixel by another manufacturer M2) is the best.

Of course, DSLR are much better than cameraphones, as can be seen on Fig. 6 (the same color scale is used).

The last figure 7 shows the dependance on the illuminant. The performance of the sensor (again the Canon EOS 400D) drops down when going from Daylight illuminant to tungsten illuminant. In particular the yellowish colors (low values on the blue channel) have the worst SNR with the neutral values. Gamut SNR depends on the illuminant through the white balance and color matrix. Since sensors are not very sensitive to short wavelengths, the blue white balance scale is usually very large. This is illustrated by the variation of Gamut SNR when switching from daylight to tungsten illuminant. There is a general loss of about 2dB. Moreover, the shape of the Gamut itself changes. The loss in red/purple ($a > 0$ and b close to 0) can be very large.

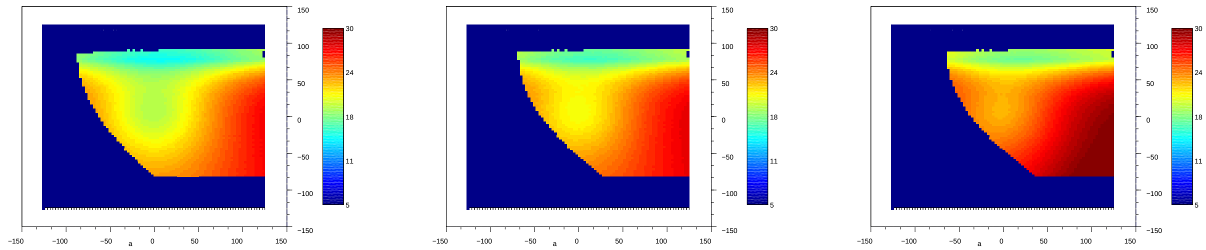

Figure 5. Gamut SNR at $L = 50$. The referent illuminant is D65. From left to right: camera module manufacturer M1, pixel pitch $2.2\mu m$, M1 with pixel pitch $1.75\mu m$, Manufacturer M2, pixel pitch $2.2\mu m$. The first two figures show that the manufacturer can have better pixel design when shrinking down the pixel size.

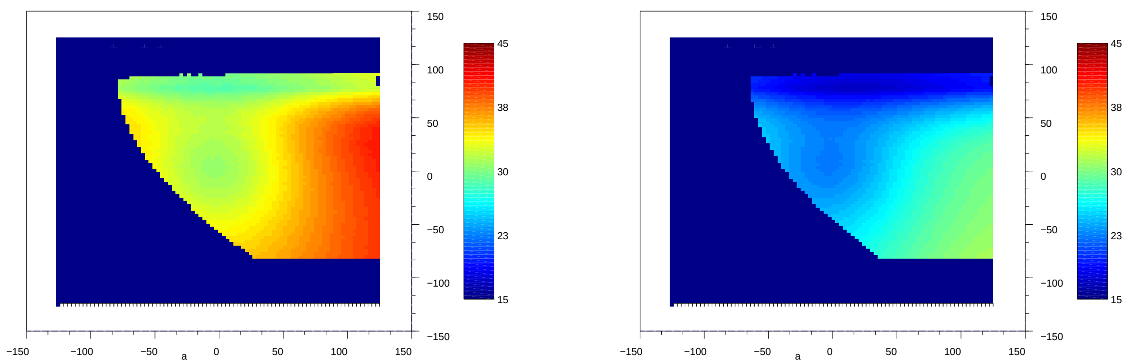

Figure 6. Gamut SNR at $L = 50$, for the Canon EOS 400D and the camera module of Manufacturer M2 (the best one in the previous plot).

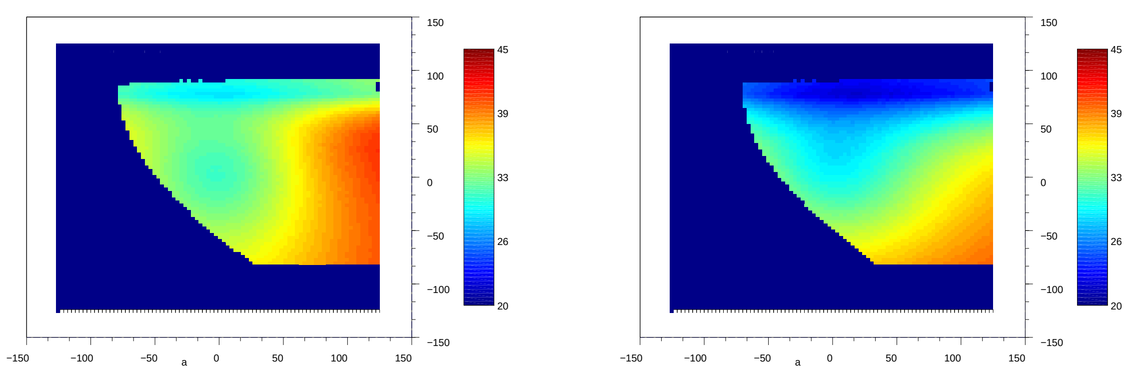

Figure 7. Comparison of Gamut SNR at $L = 50$ for Canon EOS 400D with illuminant D65 (left) and A (right). The lack of sensitivity in the blue channel, the large white balance scales are critical for the noise values after color rendering. The loss is usually about 2dB.

5. CONCLUSION

The quality of color reproduction by a camera is determined by the spectral responses of the sensor, but also by the electronic characteristics that determine the sensor noise. The notion of Gamut is not sufficient to describe this. Two measurements are proposed to take noise into account. The color sensitivity is a global measure counting the number of colors the sensor can render, up to noise. The Gamut SNR shows the distribution of noise on the sensor, after a necessary color calibration. Both notions can be used for camera raw conversion tuning, especially in low light conditions. In this case, colors are usually desaturated in order to limit noise, particularly for neutral tones. Therefore, there is a trade-off between color accuracy (ΔE error) and noise that has to be determined by experimental subjective experience. In a further work, we will present a measurement of metamerism which is a necessary complement of the color sensitivity and Gamut.

ACKNOWLEDGMENTS

The authors would like to thank Hervé Macudzinski for carrying out the experiments.

REFERENCES

1. J. Morovic and P. Morovic, "Determining color gamuts of digital cameras and scanners," *Color Research and Application* **28**(59), 2003.
2. F. M. Verdú, J. Pujol, and P. Capilla, "Calculation of the color matching functions of digital cameras from their complete spectral sensitivities," *Journal of Imaging Science and Technology* **46**(1), 2002.
3. F. M. Verdú, J. Pujol, and P. Capilla, "Characterization of a digital camera as an absolute tristimulus colorimeter," *Journal of Imaging Science and Technology* **47**(4), 2003.
4. F. M. Verdú, M. Luque, P. Capilla, and J. Pujol, "Concerning the calculation of the color gamut in a digital camera," *COLOR research and application* **31**(5), 2006.
5. B. Horn, "Exact reproduction of colors," *CVGIP* **26**, pp. 135–167, 1984.
6. G. Finlayson and M. Drew, "Constrained least-square regression in color spaces," *Journal of Electronic Imaging* **6**(4), 1997.
7. M. Vrhel and H. Trussel, "Color correction using principal components," *Color Research and Application* **17**, pp. 328–338, 1992.
8. ISO, *Graphic technology and Photography - Colour characterisation of digital still cameras (DSCs) - Part 1: Stimuli, metrology and test procedures)*, ISO, 2006.
9. J. Buzzi, F. Guichard, and H. Hornung, "From spectral sensitivities to noise characteristics," *Electronic Imaging* , 2007.
10. ISO, *Photography - Electronic still-picture cameras - determination of ISO speed*, ISO, 1998.