Influence of Spectral Sensitivity Functions on Color Demosaicing

David Alleysson[1], Sabine Süsstrunk[2], and Joanna Marguier[2]
[1]Laboratory of Psychology and NeuroCognition (LPNC)
University Pierre-Mendes France, Grenoble, France (UPMF)
[2]Audiovisual Communications Laboratory (LCAV)
Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland
Correspondence : David.Alleysson@upmf-grenoble.fr

Abstract

Color images acquired through single chip digital cameras using a color filter array (CFA) contain a mixture of luminance and opponent chromatic information that share their representation in the spatial Fourier spectrum. This mixture could result in aliasing if the bandwidths of these signals are too wide and their spectra overlap. In such a case, reconstructing three-color per pixel images without error is impossible. One way to improve the reconstruction is to have sensitivity functions that are highly correlated, reducing the bandwidth of the opponent chromatic components. However, this diminishes the ability to reproduce colors accurately as noise is amplified when converting an image to the final color encoding. In this paper, we are looking for an optimum between accurate image reconstruction through demosaicing and accurate color rendering. We design a camera simulation, first using a hyperspectral model of random color images and a demosaicing algorithm based on frequency selection. We find that there is an optimum and confirm our results using a natural hyperspectral image.

I. Introduction

Most digital cameras today use a single CCD or CMOS sensor. To capture color information, a color filter array (CFA) is placed in front of the sensor. This array, usually composed of three filters, limits the sensitivity of each photocell to a single part of the visible spectrum. Consequently, each pixel of the CFA image only contains information about this limited range, i.e. a single color response. However, three colors per pixels are needed to render images on a particular display device. An algorithm called *color demosaicing* is therefore applied to regenerate the missing colors. In a CFA image, spatial and chromatic information are mixed together in a single lattice, as only one chromatic sensitivity is available at each pixel. This mixture could result in aliasing if the lattice information capacity is not enough. For this reason, the design of a CFA camera will always involve a trade-off between spatial resolution and color accuracy.

It has been shown that spatial juxtapositions of chromatic samples (as in a CFA) could be expressed as a modulation of luminance and opponent chromatic signals. Since the arrangement of color filters in a CFA is regular, luminance and opponent chromatic signals have specific locations in the spatial Fourier domain. Moreover, luminance is defined in such a way that it has maximum spatial resolution, whereas opponent chromatic signals are sub-sampled [1]. These properties have several consequences for demosaicing. First, it is possible to design a demosaicing algorithm by selecting frequencies corresponding to the position of luminance and opponent chromatic signals in the spatial Fourier spectrum. The three-color per pixel image is reconstructed as the sum of estimated luminance and interpolated chrominance [2]. As the estimators are based on frequency selection, they can be designed as uniform linear convolution filters and offer the best compromise between accuracy and efficiency [4]. Second, since luminance is defined with maximum resolution, it is not subject to interpolation and thus is able to carry all spatial information.

However, luminance and opponent chromatic signals share the same two-dimensional Fourier space for their own representation. Artifacts may result in the demosaiced image if their representations overlap (alias). Thus, the design challenge with a CFA camera is to arrange the spatial bandwidths of luminance and opponent chromatic signals in such a way that their spectra do not overlap [4]. In particular, the bandwidth of opponent chromatic signals could be reduced by increasing the correlation between chromatic sensitivities of color filters. The underlying idea is that if the three filters are identical, the chromatic signals vanish, and then the sensor becomes a black-and-white imager. An all-pass filter estimator can exactly select the luminance information. In that case, demosaicing is perfect, but no color information is available.

Actually, the representation of luminance and chromatic opponent signal in the CFA image depends on the spatial-chromatic correlation between color planes, which can be partially controlled through the sensitivity functions of the sensor plus filter system. Unfortunately, if increasing correlation improves demosaicing, it reduces the ability to render colors in the final color encoding because noise is more amplified when the transformation becomes stronger [5].

In this paper, we discuss the influence of spectral sensitivity functions on the spatial frequency spectrum of a CFA image and judge the quality of the reconstruction using a *demosaicing algorithm by frequency selection* [2,4]. We investigate if there is an optimum between increasing the correlation of the filter sensitivities and the accuracy of color rendering. We chose to use a random hyperspectral image as input (with diagonal covariance) to prevent natural spatial and chromatic correlations. Usually, natural images have spatial and chromatic correlations that could modify the optimal parameters of the sensitivity functions. To prevent for a particular optimization based on a particular image database, we therefore chose to work with random hyperspectral images.

While the idea is intuitive, designing an optimization procedure is challenging because there are many parameters to consider: position, size and shape of the three filter sensitivities, arrangement of the filters in the CFA, and the parameters of the frequency selection algorithm. Moreover, the lack of reliable natural color image models can render the parameter estimation image dependent. In a first attempt, we study principally the feasibility, using artificial random images. To transpose this study to real capture devices, the parameters of the sensitivity functions and the point spread function of the optical system should be adapted.

The paper is organized as follow (see Figure 1 for reference). Using artificial random hyperspectral images, we show that increasing correlation between sensitivity functions increases the demosaicing quality (PSRN demosa) but decreases color rendering quality (PSNR color). Then we take into consideration the whole image processing chain and show that an optimum exists (PSNR total). Finally, we confirm this result on a natural hyperspectral image.
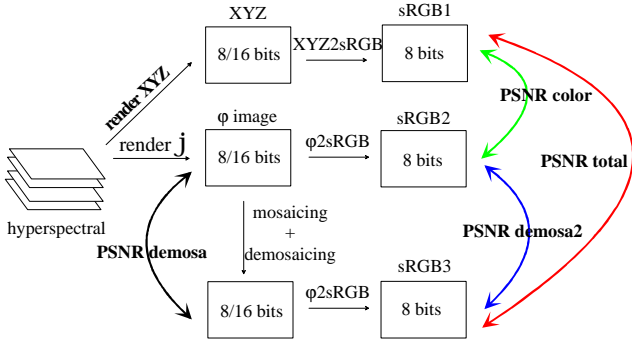

Figure 1: Synopsis of the simulations used in the paper. See text for explanation.

II. B/W versus Color Imager

Image model

Natural environments emit or reflect photons that are captured by the photosensitive elements of the camera. Even if a photon signal is a continuous function of space and wavelengths, we may consider the function of incident light $E(x,y,\boldsymbol{l})$ as a discrete three-dimensional function. Light efficiency is not constant over wavelength, and also depends on the sensitivity functions of the color filters and the sensor. Thus, the signal carried by each photocell in a digital camera can be written as:

$$C_i(x,y) = \sum_{l=1}^{N} E(x,y,\boldsymbol{l})\boldsymbol{j}_i(\boldsymbol{l}) , \ i \in [R,G,B], \qquad (1)$$

$$(x,y,\boldsymbol{l}) \in \mathbb{N}^3$$

where $\boldsymbol{j}_i$ represents the sensitivity of the system sensor plus filter of type i. In this paper, we use artificial sensitivity functions based on a Gaussian distribution in the wavelength domain. They are expressed as follow:

$$\boldsymbol{j}_i(\boldsymbol{l}) = K_i e^{-\frac{(\boldsymbol{l}-\boldsymbol{m}_i)^2}{\boldsymbol{s}_i^2}} \quad (2)$$

$$\boldsymbol{l} = [400:10:700] \ \text{and} \ \ K_i = \sum_{l=400}^{700} e^{-\frac{(\boldsymbol{l}-\boldsymbol{m}_i)^2}{\boldsymbol{s}_i^2}}$$

$\boldsymbol{m}_i$ and $\boldsymbol{s}_i$ are the parameters that determine maximum and width, respectively, of the sensitivity function of filter i. K_i is a normalization factor that sets the quantum efficiency of each sensitivity function equal to unity in the spectral interval (400 to 700 nm) used in this simulation (see Figure 2). Thus, if $E(x,y,\boldsymbol{l})$ contains 31 matrices of range 0-1, the corresponding RGB image computed through the $\boldsymbol{j}_i$ sensitivity functions keep the same range.
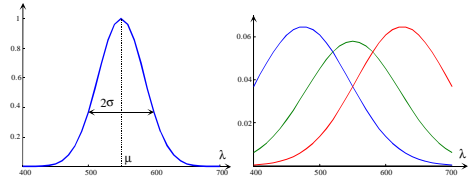

Figure 2: Representation of simulated sensitivity functions. (a) $\boldsymbol{m}$=550, $\boldsymbol{s}$=50: the width of $2\boldsymbol{s}_i$ is located at approximately 37% of the curve height. (b) The normalization factor guarantees a unitary integral of each sensitivity function, even if the curve is cut due to range limitation. $\boldsymbol{m}_i$=(475, 550, 625), $\boldsymbol{s}_i$=100.

For a more realistic simulation, we also take into account the optics of the camera. The function $E(x,y,\boldsymbol{l})$ that describes incident light on the sensor corresponds to an illumination function convolved with a low pass filter, which represents the effect of the optical system's point-spread function (PSF). We model the PSF by a Gaussian function. To simplify the computation, we apply a convolution in spatial Fourier domain. If $\hat{F}$ and $\hat{E}$ represent the Fourier transform of scene radiation and radiation after passing through the optics, respectively, we have:

$$\hat{E}(f_x,f_y,\boldsymbol{l}) = \hat{F}(f_x,f_y,\boldsymbol{l})e^{-\frac{f_x^2+f_y^2}{s^2}} \qquad (3)$$

Where $F(x,y,\boldsymbol{l})$ is a tri-dimensional random matrix. Increasing parameter s reduces the blurring effect symmetrically for x and y spatial dimensions, and vice versa.

2

With this model, generating an RGB image requires five parameters: three for the maximum of wavelength sensitivity for the three filters, one for the variance (we choose the three variances to be equal) and one for the optical blur.

Correlation is optimal for demosaicing

In a CFA image, each pixel contains only one color response as opposed to three color responses in regular RGB color images. It is possible to express the whole signal of a CFA image in terms of C_i by taking into account the spatial arrangement of the color sensitivities on the sensor. As already shown in [1], the whole CFA signal is expressed as:

$$I_{CFA}(x,y) = \sum_i C_i(x,y)m_i(x,y) \qquad (4)$$

where $m_i(x,y)$ are sub-sampling functions taking values 1 or 0, depending on the presence or absence of sensitivity i at position (x,y).

The spatial Fourier transform of a CFA image following the x and y variables is given by:

$$\hat{I}_{CFA}(f_x,f_y) = \sum_i \hat{C}_i(f_x,f_y) * \hat{m}_i(f_x,f_y) \quad (5)$$

where $\hat{}$ represent the spatial Fourier transform and $*$ the convolution operator. f_x and f_y are the spatial frequencies corresponding to x and y. Since $m_i(x,y)$ functions are periodic, the global Fourier transform of a CFA image is composed of distinct energy regions.

In Figure 3(d), we can distinguish nine regions where energy is concentrated. The one in the center corresponds to luminance, and the eight in the corners and centers of each side correspond to opponent chromatic signals. In this case, the optical blur is designed to guarantee that the nine regions are well separable. In Figure 3(e), the optical blur is less strong, resulting in an overlap between the different regions that can generate artifacts in demosaicing.
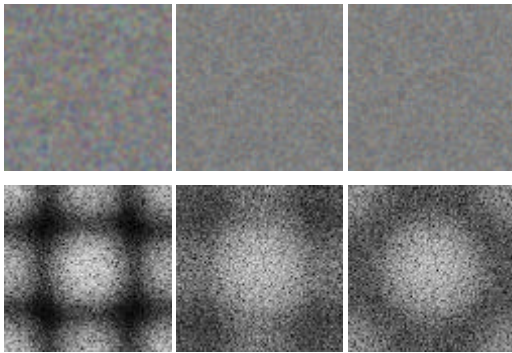

*Figure 3: Random images (128x128), sampled with sensitivity functions with **m**=(475, 550, 625) and their spatial Fourier spectrums (a) RGB image with s=30, **s**$_i$=50; (b) RGB image with s=50, **s**$_i$=100; (c) same image as (b); (d) Fourier spectrum of (a) after sub-sampling according to the Bayer CFA; (e) idem of (b); (f) idem of (c) but with a CFA following the Bayer spatial arrangement but having twice more blue than red and green pixels.*

As already proven elsewhere [2], the spatial arrangement of filters proposed by Bayer [3] is the best spatial arrangement of three colors on a square grid. From the point of view of demosaicing, however, it is better to have a CFA with twice more blue than red and green pixels [4]. The resulting arrangements of the opponent chromatic spatial frequencies will reduce aliasing in the reconstruction. This is illustrated in the Figures 3(e) and 3(f), as well as Figures 4(b) and 4(c).
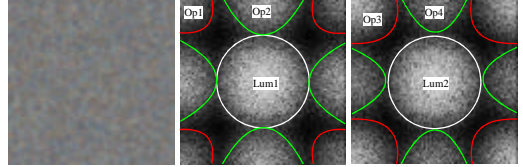

Figure 4: Spatial Fourier representation of a CFA image. (a) Original image; (b) spatial Fourier spectrum of its Bayer CFA image; (c) same of a Bayer CFA with twice more blue than red and green pixels.

$$Lum1 = \hat{C}_R + 2\hat{C}_G + \hat{C}_B, \quad Op1 = \hat{C}_R - 2\hat{C}_G + \hat{C}_B, \quad Op2 = \hat{C}_R - \hat{C}_B,$$

$$Lum2 = \hat{C}_R + \hat{C}_G + 2\hat{C}_B, \quad Op3 = \hat{C}_R + \hat{C}_G - 2\hat{C}_B, \quad Op4 = \hat{C}_R - \hat{C}_G$$

In Figure 4, we additionally indicate the content of the opponent chromatic channels. From the same image (Figure 4(a)) we compute its spatial Fourier transform after sub-sampling according to the Bayer CFA (Figure 4(b)) and a "Blue-Bayer" CFA with twice more blue than red and green pixels (Figure 4(c)). In both spatial Fourier representations of Figures 4(a) and 4(b), luminance occupies the same area, even if it is not defined by exactly the same amount of R, G and B. However, chrominance signals have different bandwidths. As can be seen in Figure 4(b), the center of the sides is composed of the Fourier spectrum of R minus B. It has a larger bandwidth than the center of the sides of Figure 4(c), which is composed of the Fourier spectrum of R minus G. Considering that the overlap between the luminance and opponent chromatic spectra is responsible for aliasing artifacts during reconstruction, it is better to have the larger chromatic spectrum located further from the luminance spectrum. That can be achieved by changing the color arrangements on the Bayer CFA, i.e. by allowing for two times more blue than red and green pixels. Note that for the simulations in this paper, we used such a CFA.

Optimal convolution filters for demosaicing

In a CFA image, luminance can be estimated by a convolution filter [1-2, 4]. We can design this filter by Gaussian functions that suppress the opponent chromatic channels and conserve as much as possible the luminance spectrum. For example, we decided to parameterize the filter as follow:

$$\hat{F}_{lum} = 1 - \sum_{c_1}\sum_{c_2} e^{-\frac{(f_x-c_1)^2+(f_y-c_2)^2}{r_1^2}} - \sum_{c_3}\sum_{c_4} e^{-\frac{(f_x-c_3)^2+(f_y-c_4)^2}{r_2^2}} \quad (10)$$

$$c_1 \in \left[-\tfrac{1}{2},\tfrac{1}{2},\tfrac{1}{2},-\tfrac{1}{2}\right], c_2 \in \left[-\tfrac{1}{2},-\tfrac{1}{2},\tfrac{1}{2},\tfrac{1}{2}\right], c_3 \in \left[-\tfrac{1}{2},0,0,\tfrac{1}{2}\right], c_4 \in \left[0,-\tfrac{1}{2},\tfrac{1}{2},0\right]$$

where c_1 and c_2 represent the centers of the opponent chromatic channels in the corner, and c_3 and c_4 the centers of each side of the Fourier spectrum. r_1 and r_2 are free parameters corresponding to the variance of Gaussian functions in the corner or center of the sides of the Fourier spectrum, as illustrated in Figure 5.
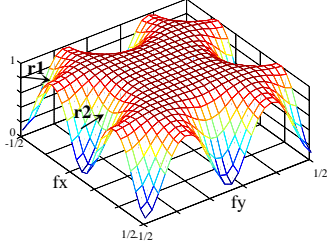

Figure 5: Representation of a luminance estimation filter in the spatial frequency domain. Parameters r1 and r2 influence the amount of the chrominance retained in the corners and centers of sides, respectively.

Once this filter has estimated luminance, the chrominance is retrieved with its orthogonal filter. This can be achieved by subtracting the estimated luminance from the CFA image. We then interpolate the opponent chromatic channels, and reconstruct the three-color by pixel image as the sum of estimated luminance and interpolated opponent colors. See [2] for details on the algorithm.

For a particular CFA image, we optimize parameters to result in maximum color peak signal-to-noise ratio (CPSNR) between the reconstructed image through the luminance estimated by the filter and the original RGB image.

Simulation

We created a random hyperspectral image $F(x, y, l)$, on which we applied a blur (eq. 2) to obtain $E(x, y, l)$. We then construct three filter sensitivity functions j_i and sample (eq. 1) the hyperspectral image to obtain a RGB image. To create a CFA image, we "mosaic" the RGB image by selecting the pixel values corresponding to a "Blue-Bayer" CFA. We then apply the frequency selection demosaicing algorithm, always computing the optimal parameters for the luminance filter. We then compute the CPSNR (called "PSNR demosa" on Figure 1) between the color image and the reconstructed image.

The first simulation we performed shows the influence of the position of the maximum sensitivity, m_i, on the demosaicing algorithm's performance. We fixed m_R and m_B to 475 and 625nm, respectively, and vary m_G from 400 to 700nm. Figure 6 illustrates clearly that the position of the sensitivity functions change the quality of the demosaicing algorithm in terms of CPSNR. If the sensitivity functions are broadband, there is only one global maximum (Figure 6(a)). If the sensitivity functions are more narrow band, there are two local maxima (Figure 6(b)).

The reason why the maximum CPSNR occurs when m_G is close to m_R can be explained by analyzing the Fourier representation. The area of the spectrum most subject to aliasing is given by the center of the sides, thus the spectrum of R-G (see Figure 4(c)). Thus, when R is equal to G, this part is removed. However, if R is equal to G, the bandwidth of the second opponent signal will increase. The optimal is therefore a G close to R, but not equal. Note that the L, M, and S cone sensitivity functions of the human visual system have similar properties; they are broadband and the maximum sensitivity of M is close to L.
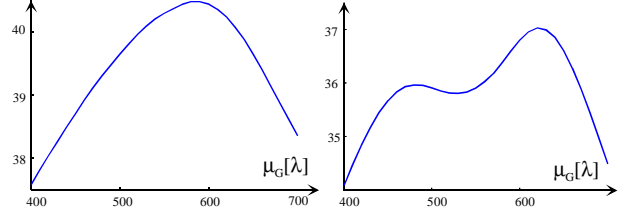

Figure 6: CPSNR between original and reconstructed image, varying parameter m_G with m_R=625, m_B=475, s=50, (a) s_i=100; (b) s_i=50.

The second parameter we tested is the bandwidth in wavelength domain of the sensitivity function, given by s_i. We constrain s_i to be equal for each of the three color sensitivity functions. Figure 7(a) shows an example of CPSNR value against s_i.
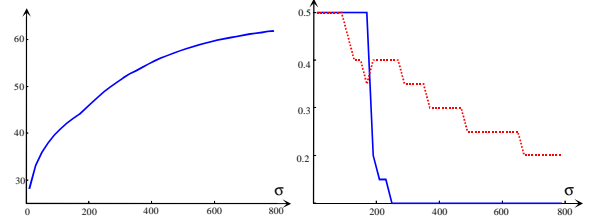

Figure 7: Influence of s_i on the demosaicing performance. (a) "PSNR demosa" for m=(475, 550, 625); (b) Evolution of radius r1(continuous line) and r2 (dotted line).

As previously mentioned, larger s_i increase the overlap between the different filter sensitivity functions, which in the extreme will result in a black and white imager. The demosaicing CPSNR consequently also increases with increasing s_i. Moreover, as shown in Figure 7(b), the parameters of the luminance filter move to an all pass filter.

Correlation is worst for color rendering

We have shown in the previous section that increasing the correlation between the sensitivity functions results in better demosaicing CPSNR. This result is antagonistic with the ability to render color. As shown, for example by Baer and Holm [5], the linear transformation from sensor encoding to final display color encoding has a large influence on noise. In this paper, we use sRGB [6] as the final display color encoding. We compare the rendering of a hyperspectral image into sRGB with an image sampled by our j_i sensitivity functions and then converted into sRGB. As the sRGB color matching functions (CMF) are based on XYZ, we consider the XYZ CMFs as the hypothetical sensor responses for the first case. The resulting image is

called sRGB1 in Figure 1. For the simulation of the second image (sRGB2), we sample the hyperspectral image with the j_i functions and then compute the best linear transformation between j_i and XYZ using mean square error. We then render the image into sRGB. The comparison of sRGB1 and sRGB2 corresponds to "PSNR color" in Figure 1.

It should be noted that in our simulation, the effect of noise is not visible if we compute in double precision. Additionally, the image model we have proposed does not contain any noise. To take into consideration the noise amplification at the color conversion step, we therefore quantize the images into either 8 or 16 bits per channel before converting into sRGB. This is analog to the A/D conversion in digital cameras.

Simulation

As the sRGB color matching functions are the basis of our reference color encoding, the closer the j_i resemble these function, the better the PSNR color will be. To confirm this hypothesis, we choose $m_R = 600$ and $m_B = 450$, thus equal to the maximum of the sRGB color matching functions, and let m_G evolve from 400 to 700 nm. Note that while we indicate the blur parameter s we have used in the simulation, it has no influence on the result because we compare here three-color by pixel images.
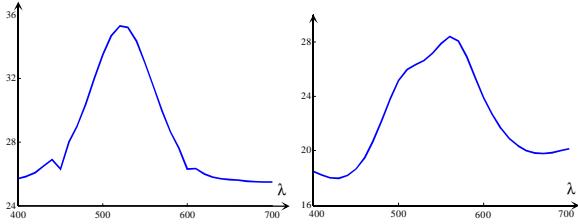

Figure 8: PSNR color between an image rendered directly into sRGB and an image rendered with j_i and then converted to sRGB. (a) m_R=600, m_B=450, s=50, s_i=50; (b) same for m_R=625, m_B=475.

The best PSNR is obtained with $m_G = 520$ (see Figure 8 (a)), which is close to the maximum of 540 nm of the sG CMF. The difference between them could arise from the quantization of the wavelength variables, or because of the different shapes of the j_i functions compared to the shape of sRGB functions. For another example, we chose $m_R = 625$ and $m_B = 475$. These values ensure uniform coverage of the visible spectrum but result in a bad PSNR. In that case, $m_G = 560$ is the optimum value (Figure 8(b)). Here we can also see a difference between the value of 550 nm that would cover the whole spectrum and the optimal value. It should be noted that for the s_i values we have used here, the noise due to quantization has no influence, even for 8 bits per channel.

For the rest of this paper, we will use the optimal maxima $m_i = (450, 520, 600)$. We can now investigate the influence of s_i on the color reproduction.
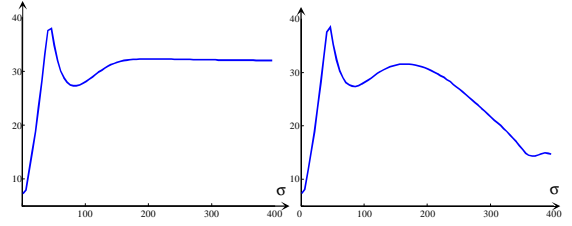

Figure 9: (a) PSNR color as a function of s_i, 16 bits/channel quantization; (b) same for 8 bits/channel quantization.

As shown in Figure 9(a), the optimal PSNR is around $s_i = 50$, independent of the quantization. The influence of quantization noise is visible at s_i larger than 150. In our example, increasing the bandwidth does not decrease the color quality for 16 bits/channel. As expected, however, the quality of the color rendering will decrease with increasing s_i when the images are quantized to 8 bits/channel (Figure 9(b)).

III Tracking optimal parameters

We have seen that increasing the correlation between the sensitivity functions improves the demosaicing quality but reduces the color reproduction quality. To find a global optimum for the sensitivity functions, we designed the following simulation.

We generate a hyperspectral random image with a given optical blur factor. We choose the m_i values as defined in the previous section, and let s_i evolve. To measure the quality of the reconstruction, we compute the "PSNR total" between sRGB1 and sRGB3, as defined in Figure 1. sRGB1 is the image rendered directly into sRGB using XYZ color matching functions as sensors. sRGB3 is an image that was sampled with the sensitivity functions j_i, quantized, mosaiced and then demosaiced, and then rendered into sRGB. We also plot "PSNR color" and "PSNR demosa" as defined in the previous section. We also compute the PSNR between sRGB2 and sRGB3 (PSNR demosa2) to take into account the demosaicing effect after rendering into sRGB.

The result is illustrated in the Figure 10. We see clearly in Figure 10(a) that there is an optimum around $s_i = 50$ for the PSNR total. For all s_i, PSNR total is lower than PSNR color and PSNR demosa2. However, there are several things to highlight. First, in contrary to PSNR demosa that increases continuously with s_i, PSNR demosa2 that takes into account the rendering to sRGB shows an increase followed by a decrease around $s_i = 80$. This is because the noise imposed by the demosaicing process is amplified by the color encoding conversion. This noise has an effect at a s_i width smaller than the s_i width found for color correction only, as can be seen by comparing with Figure 9(b), where the PSNR color starts to decrease around $s_i = 150$.

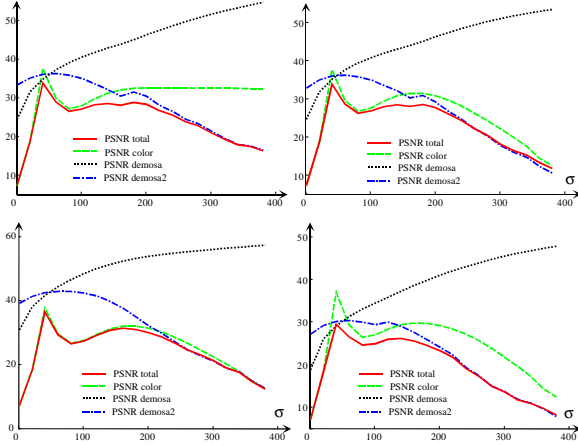*Figure 10: (a) PSNR function of s_i with s=50, $\mathbf{m}$=(450, 520, 600), 16 bits (b) same with 8 bits (c) same with s=30, (d) same with s=80.*

Figure 10(c) and 10(d) are the PSNR results for optical blur values s of 30 and 80, repectively. These figures show that demosaicing results vary for these values, starting with a higher value when s decreases. But PSNR color and consequently PSNR total are not modified, as long as the PSNR demosa is not too small (Figure 9(d)).

In Figure 11, we show an example of an image rendered directly (sRGB1) and an image rendered through the whole chain (sRGB3), as well as the optimal sensitivity functions with $s = 44$.
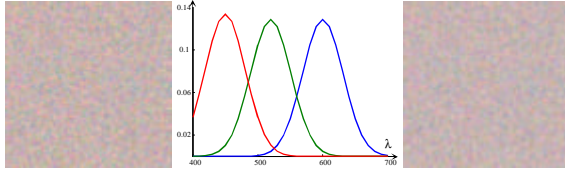

Figure 11: (a) sRGB1 (b) optimal sensitivities function (c) sRGB3 using optimal parameters

In order to confirm our result on natural images, we have tested our framework on a particular hyperspectral image [7]. A white balancing correction is performed by normalizing the white square of the MacBeth color checker in the image, as our simulation does not take illuminant variations into account. The image is also cropped, and the simulation is executed on a 128x128 window.
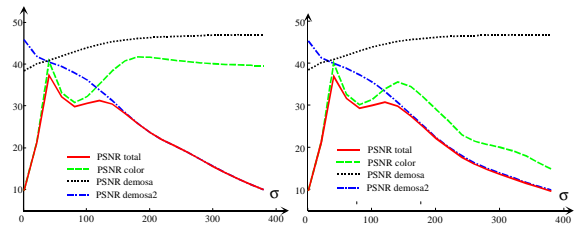

Figure 12: Simulation on a real hyperspectral image. (a) for 16 bits and (b) for 8 bits.

As shown in figure 12, we found the same kind of behavior for the four different PSNR, meaning that our image model is close to real images. Nevertheless, the optimal s_i value is a bit smaller. Note also that the PSNR values are higher than for a random image. That is particularly true for "PSNR demosa" which shows an initial value of around 40dB. This value was closer to 25dB in random images. This is certainly due to the natural correlation in space and wavelength inherent to natural images.

IV Conclusion

The goal of this paper is to find the optimal sensitivity functions for CFA filters by examining the trade-off of increasing their correlation to improve demosaicing algorithm performance and decreasing them to improve their color rendering abilities. As shown in our simulation on random hyperspectral images, we can find an optimum for every case.

However, the methods presented in this paper are a simulation, limited to sRGB as end color encoding. They also do not take into account white-balancing and other color corrections. Additionally, using a more appropriate quality metric than PSNR could be discussed. Finally, we limit our simulation to equal quantum efficiency for all sensors and do not consider different illuminant spectral power distributions. For real applications, the shape of the sensitivity functions, optical blur, display color space, capture noise, etc. should be take into consideration.

Finally, this paper confirms that demosaicing by frequency selection, or equivalently using linear convolution filters, is a fast and accurate way to reconstruct three-color per pixel images. We have also shown that color conversion and quantization has a larger influence on quality compared to the demosaicing algorithm. Also, we confirm that good camera design is a compromise between spatial blurring by the optics, the shape of the sensitivity functions, the ability to reconstruct good spatial acuity, and color accuracy.

References

[1] D.Alleysson and J.Hérault, "Interpolation d'images couleurs sous échantillonnées par un modèle de perception," *Proc. GRETSI'2001*, Toulouse, Sept. 2001.

[2] D.Alleysson, S. Süsstrunk and J. Hérault, "Color Domosaicing by estimating luminance and opponent chromatic signals in the Fourier domain," *Proc. CIC10*, 2002.

[3] B.E. Bayer, 1976, Color imaging array, US Patent 3,971,065

[4] D. Alleysson, S. Süsstrunk and J. Hérault, "Accurate color demosaicing inspired by the human visual system," submitted to *IEEE Trans. Image Processing*, March 2003.

[5] R.L Baer and J. Holm, "A model for calculating the potential ISO speeds of digital still cameras based upon CCD Characteristics," *Proc. PICS*, 1999.

[6] IEC 61966-2-1 (1999): Colour management – Default RGB colour space – sRGB,.

[7] D. H. Brainard, Hyperspectral Image Data, http://color.psych.ucsb.edu//hyperspectral/.