
From Photon to Pixel

The Digital Camera Handbook

Henri Maître

Color Section

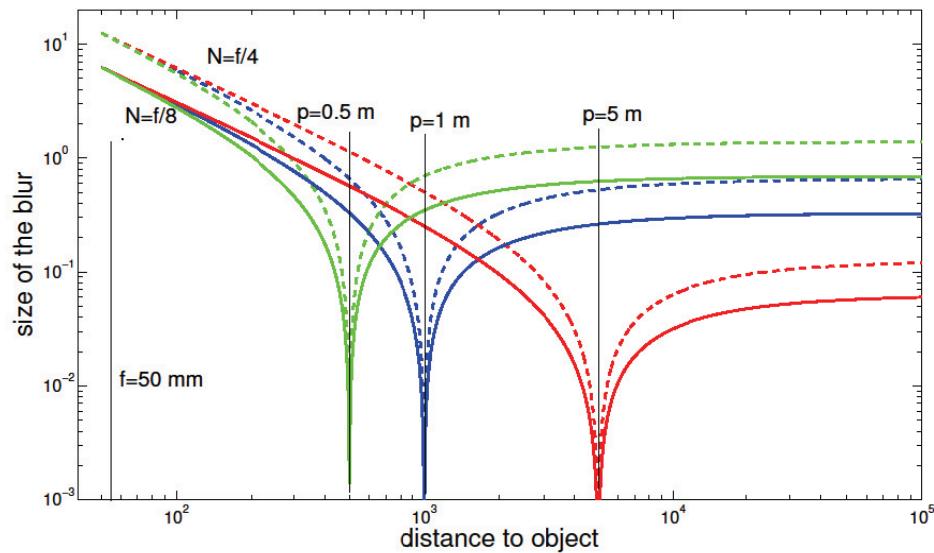


Figure 2.2. Evolution of the size of the blur spot relatively to the displacement along the optical axis. The scales on the two axes are in millimeters. The focal length is fixed at 50 mm. Three focusing distances are studied: 0.5, 1 and 5 m, as well as two apertures of the diaphragm: $N = f/4$ and $N = f/8$. On the left, in a log-log plot of the focal point F at a distance of 100 m, On the right, in a semi-log plot

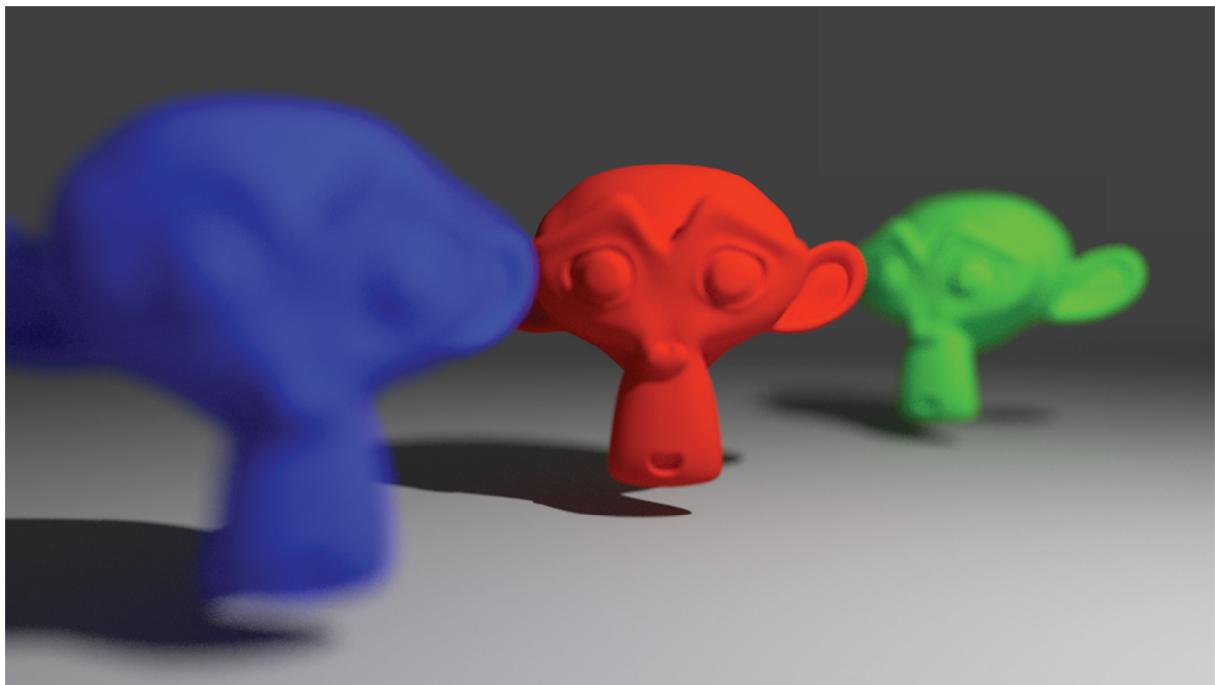


Figure 2.4. The depth of field blur when creating a synthesis image is simulated by a partial occultation effect by sending rays through a partially masked objective (picture T. Boubekeur)

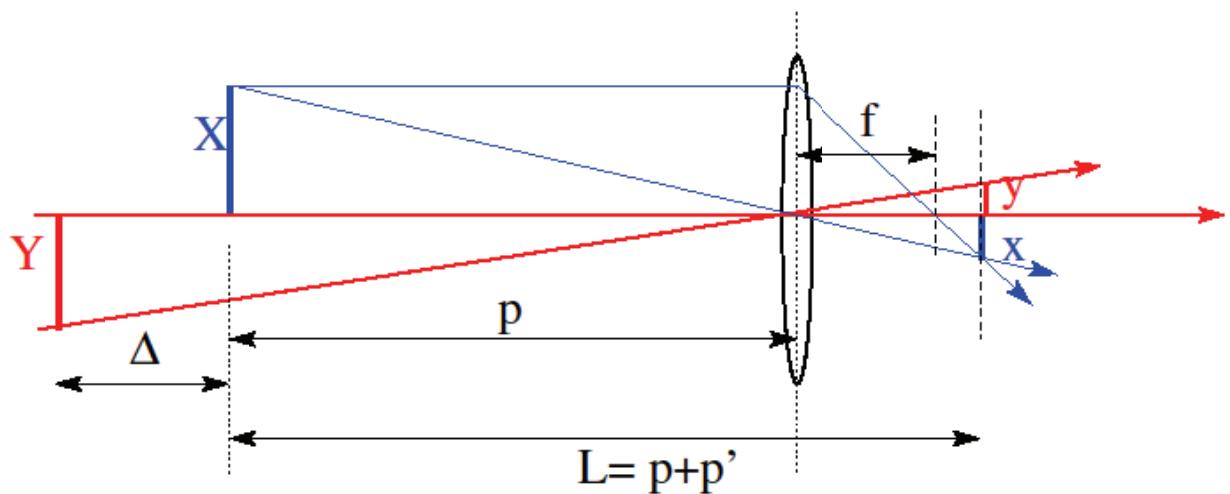


Figure 2.16. On the condition that the depth of field conditions be verified, an (almost) sharp image of an object Y (of the same size as X) can be obtained after focusing on the object X. The image y of Y appears then (in this case where the object has moved backward) smaller than the image of X

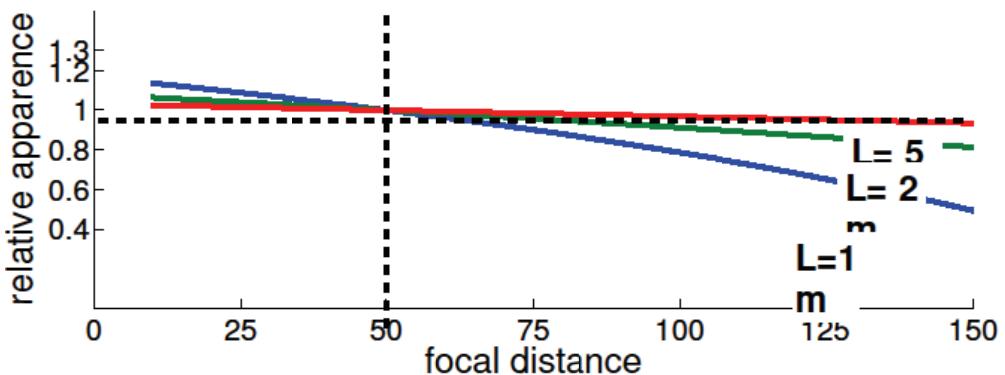


Figure 2.17. Impact of the focal length on the relative apparent size of two objects with the same size, one placed at a distance L , the other placed at a distance $3L$ from the photographer, for an objective whose focal length varies from 10 to 150 mm (see Figure 2.16). The relative size seen by a 50 mm lens is taken as the reference. The three curves correspond to values of L of 1, 2 and 5 m

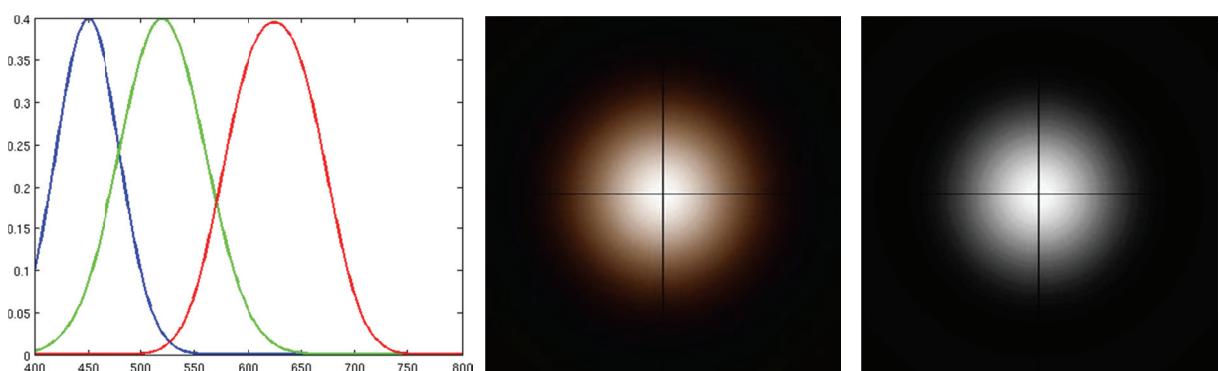


Figure 2.27. PSF obtained with white light. The sensor has three sensitive CCDs according to the curves of the figure on the left. The light spectrum is uniform on the visible band from 0.4 to 0.8 μm . In the center, the PSF obtained by summation of the contributions of all wavelengths: the spot is dominated by the red wavelengths outside the axis. On the right, hypothesis of an identical average diffraction for all the wavelengths. The spot would be monochrome

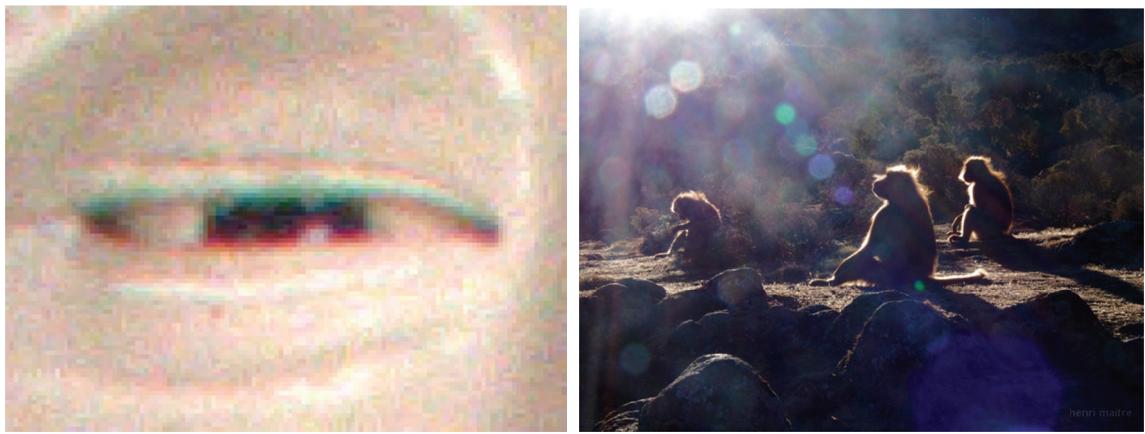


Figure 2.30. On the left, chromatic aberrations. The red or blue edging that affects the contours above and below the eye can clearly be seen (the full image where this detail is taken from is in Figure 8.2). Right image: multiple reflections on the various lenses of the objective: the heptagonal shape of these reflections comes from the diaphragm composed of seven plates opening like a hand fan. The shape of the diaphragm is involved in these reflections, but also in the shape of the blurring of the objects that are focused

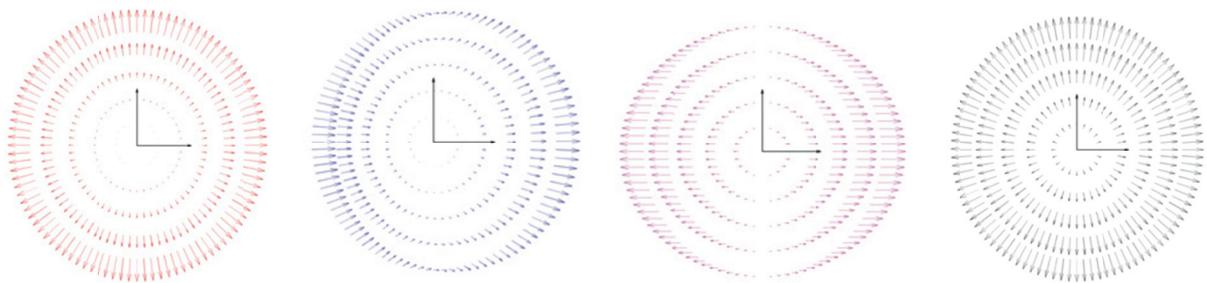


Figure 2.31. Spherical, coma, astigmatism and field curvature aberration, according to [TAN 13]. The graphics represent the deformation of the image point according to its position in the field

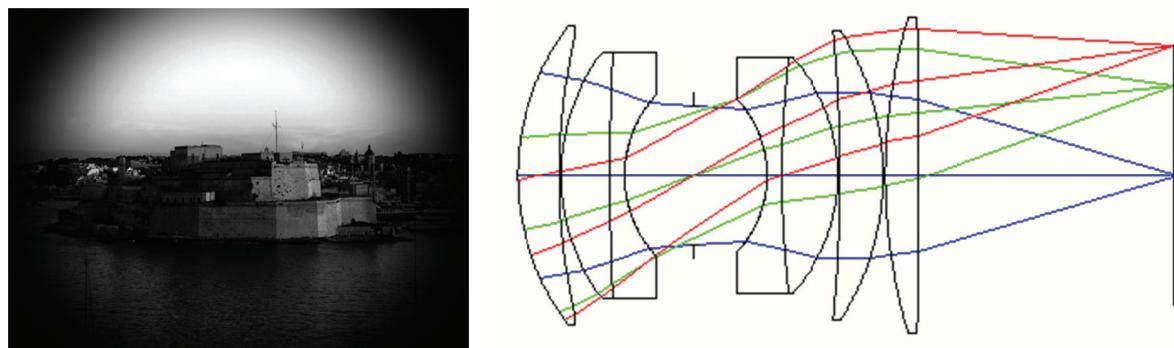


Figure 2.32. On the left: vignetting effect on a wide-angle image. On the right: tracing of the optical paths in an objective. The red radius corresponds to an image point of the edge of the field. Only the rays traveling from the lower half of the objective will cross the diaphragm in the center of the objective (©Radiant Zemax)

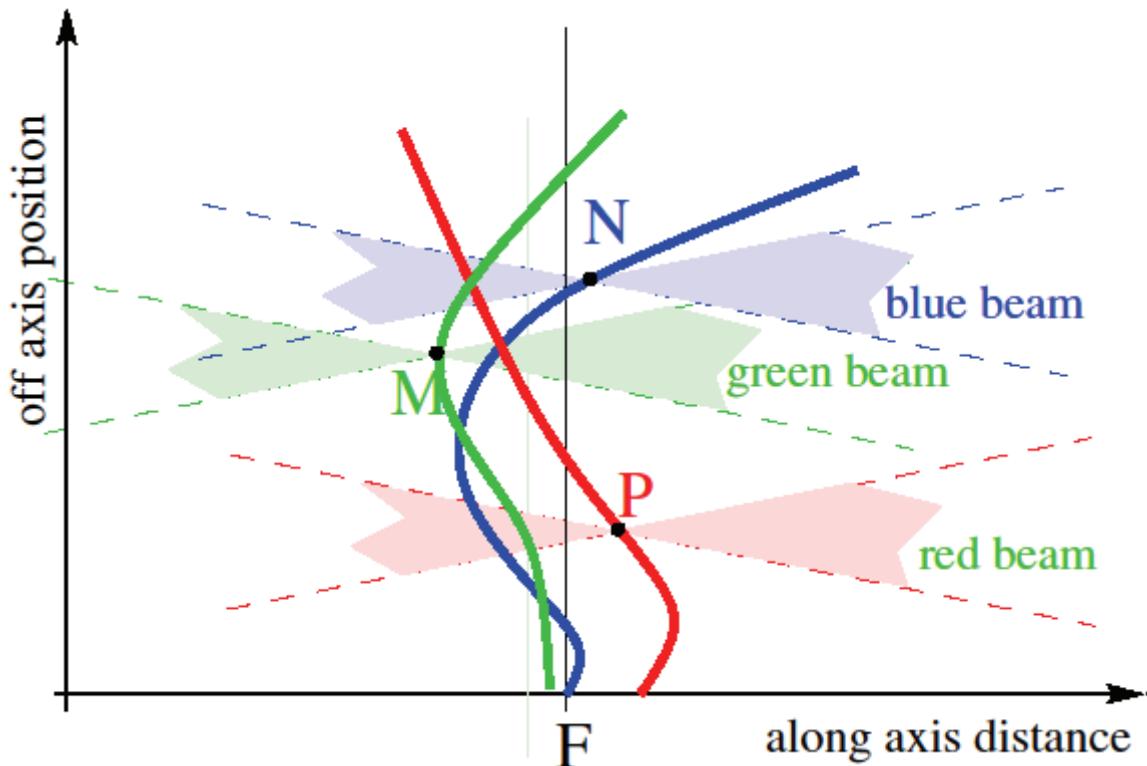


Figure 2.34. Example of the imperfect result of a correction of aberrations. Complex compromises are made without, however, completely canceling the defect. Here, the blue is too corrected at the edge of the field, while the red is not sufficiently corrected on the axis. Three beams red, green and blue focus on three points M, N and P on the corrected focal nappes. It is very difficult to determine the real focal plane F of such a system

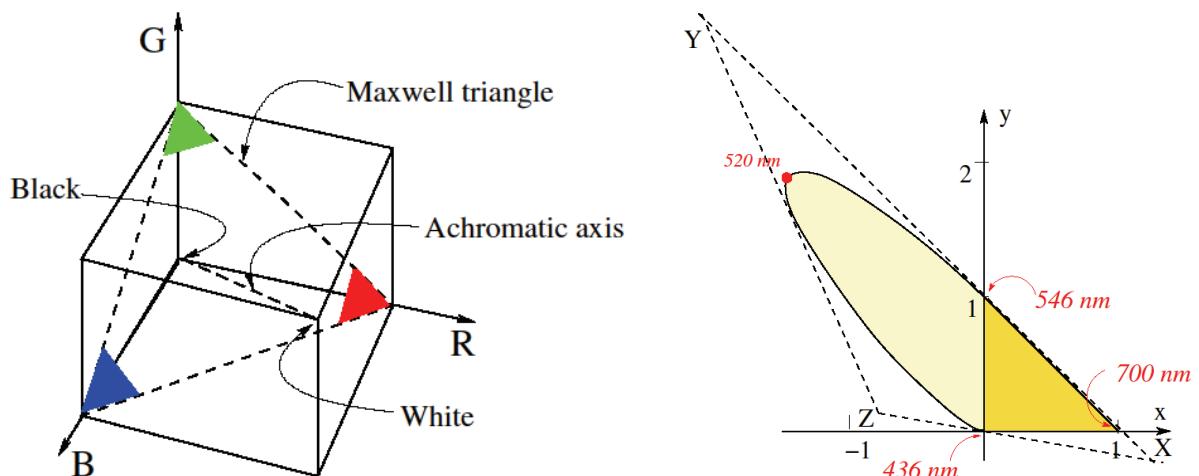


Figure 5.3. On the left, the cube of RGB colors and the Maxwell triangle. On the right, the chromatic diagram (x, y) built on constant brightness RGB ($R+G+B = K$): the origin is placed in B and only the normalized R and G coordinates are retained. The site of pure frequencies involves a very important part of negative abscissa. The triangle XYZ is the one that has been selected to define the axes of the CIE 1931 observer in order for any real color be obtained by a positive combination of three terms

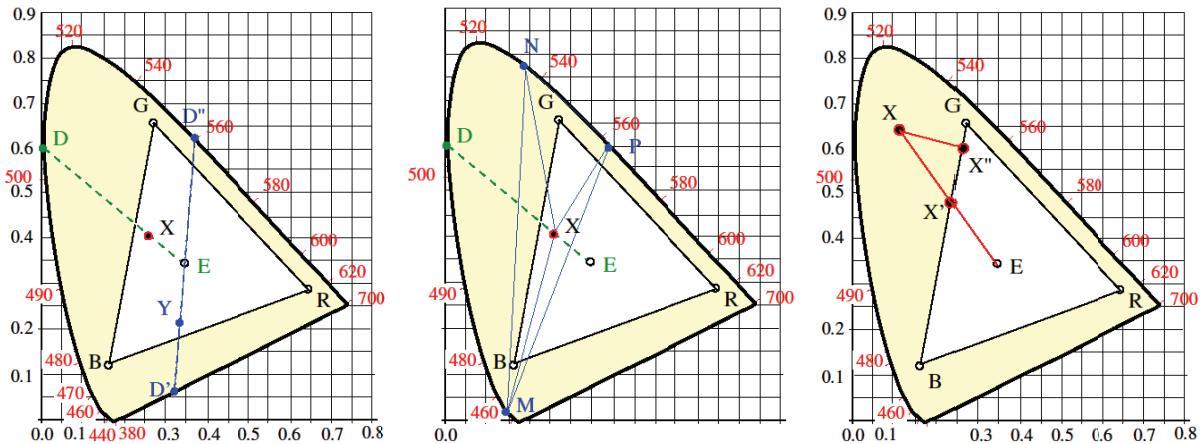


Figure 5.5. Color space with three RGB primaries and a reference white noted E . Left: point X has for dominant wavelength a blue-green one at 505 nm (point D). The point Y , on the other hand, has no dominant wavelength (the point D' corresponds to a purple which is not a pure radiation). By convention, the wavelength of the complementary (point D'') at 562 nm is attributed thereto. In the center: example of metamerism. The point X , single color point in RGB representation, could have been created either by a mixture of blue-green and white (linear interpolation on the line DE), or by a mixture of blue (at 460 nm), of green (at 532 nm) and red (at 564 nm) (barycentric interpolation in the triangle MNP), or even by an infinite number of other combinations of different radiations. They will all be seen as a same color X by the observer. Right: the point X is out of the gamut. It can be represented in the RGB space by point X'' (the nearest to X) or by point X' (same hue as X , but desaturated)



Figure 5.6. The color image is decomposed into its chromatic channels: red, green and blue which are represented here in grayscale



Figure 5.7. The color image of Figure 5.6 is decomposed here into its chromatic components. On the left, the hue is represented in hues: that is the pure wavelengths from yellow to red. To its right, the hue is represented in grayscale by an angle comprised between O (represented in black, but therefore coding the yellow) through the green (absent from this image and the blue which appears thus in gray) and 2π (represented in white, but coding the red color). On the right, luminance image (this is often the one that is chosen to transform a color image into a black and white image). On its left, the representation of saturation (faded hues appear in black)



Figure 5.9. White balance. From left to right: original image, whose parameters (illuminants, camera settings) are unknown, as recorded on chip (during viewing the primaries are assumed to be those of the sRGB space); white balance by the von Kries model, under the gray-world hypothesis, then under the white-page hypothesis, finally on the right, correction made with the (unknown) algorithm of the camera. The very wide variety of renderings can be observed in this difficult situation where the light sources are artificial and the colors strong



Figure 5.11. Decomposition of an image in its three RGB components. The original image is on the top left. The red channel on its right is represented in grayscale for an easier comparison with the green and the blue channels in Figure 5.12



Figure 5.12. Decomposition of an image into its three RGB components, the green and blue channels (continued from Figure 5.11)

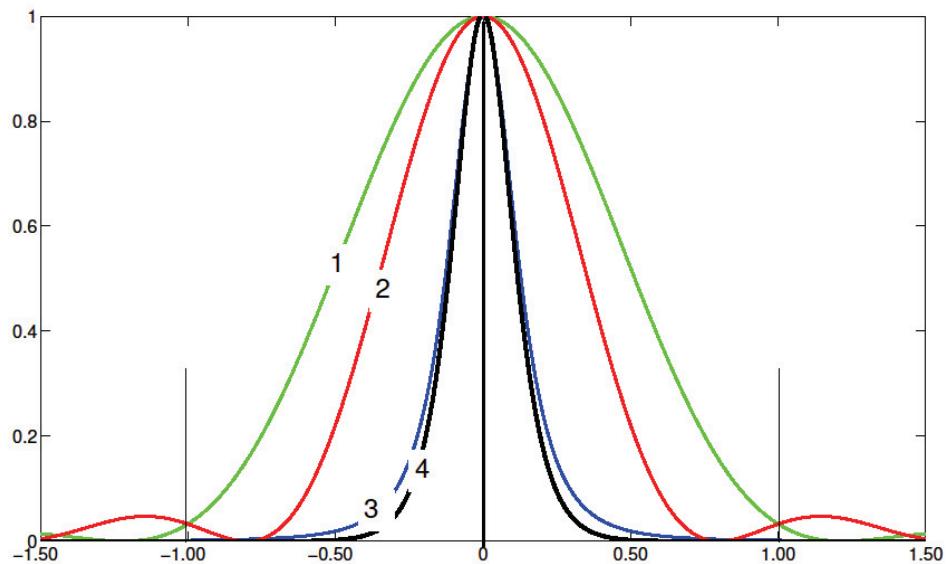


Figure 5.14. The spectrum of the image signal (curve 1), maximal at frequency zero, the filter imposed by the diaphragm (curve 2) and that imposed by the integration on the sensor (curve 3). If the period of the sensor defines useful frequencies in the band $u \in [-1, 1]$ we have arbitrarily chosen a diaphragm that cancels out at ± 1.22 and a sensitive cell such as $\mu = 1.25\eta$. The signal resulting from the two filterings prior to periodization by the sensor matrix is represented by the curve 4

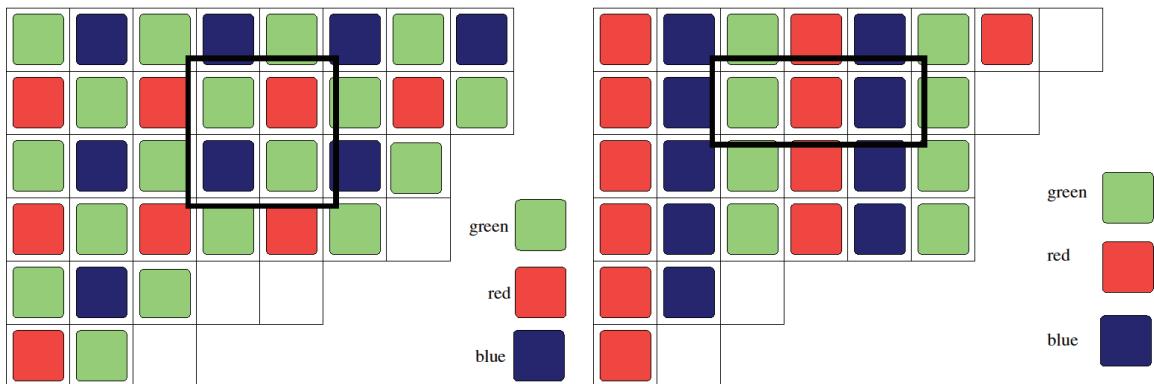


Figure 5.17. On the left, Bayer filter array; half the sites are sensitive to the green component, a quarter to the blue, a quarter to the red. The repetition element step of the pattern is 2×2 . On the right, a chromatic stripe selection array. A third of the pixels is assigned to each color. One direction is heavily under-sampled while the other is correctly sampled. The pattern is 3×1 . The configurations of the basic patterns of the periodicity are traced in black

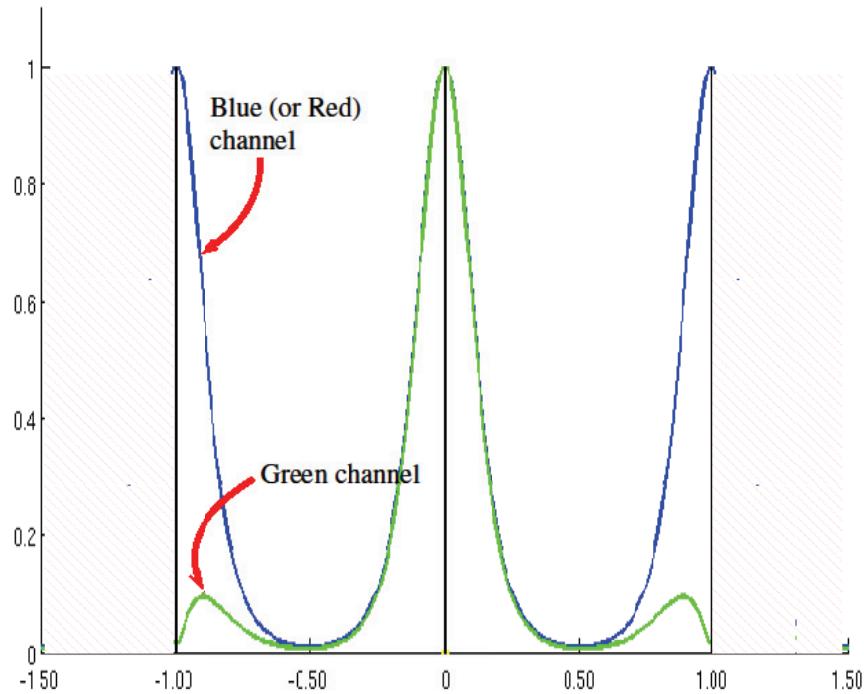


Figure 5.18. Given a dimension, the power spectral density of an image created by a sensor using a Bayer mosaic filter: green channel on one hand, blue channel (or red) on the other hand. The periodization of the photosites corresponds to a bandwidth between -0.5 and $+0.5$. The blue channel, sampled every two pixels, undergoes a very strong aliasing that the green channel barely exhibits

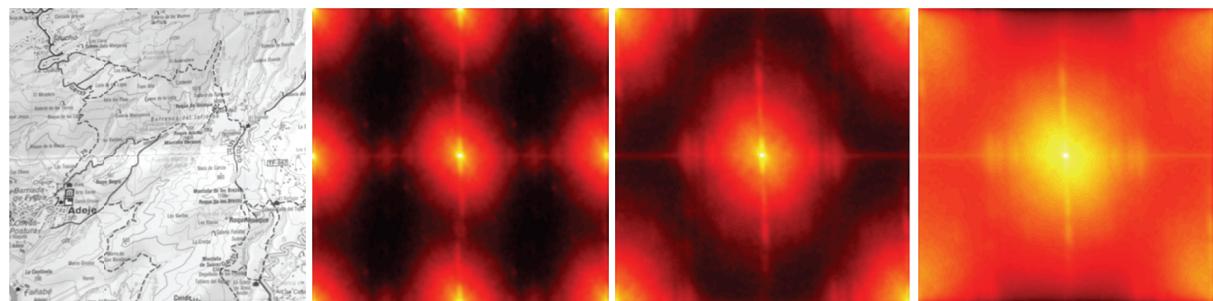


Figure 5.19. On the left, the green channel, reconstructed by interpolation, of a RAW image. The three images on the right represent the frequency spectrum module of the blue channel before interpolation, then of the green channel before interpolation, and finally of the green channel after interpolation. To facilitate the readability of these spectra, they have been filtered by a low-pass Gaussian filter, and represented in logarithmic scale. It can be emphasized that aliasing is twice stronger from the blue channel compared to the green channel and that the role of the interpolation in rejecting noise orders of the green channel outside the bandwidth thus reduce aliasing

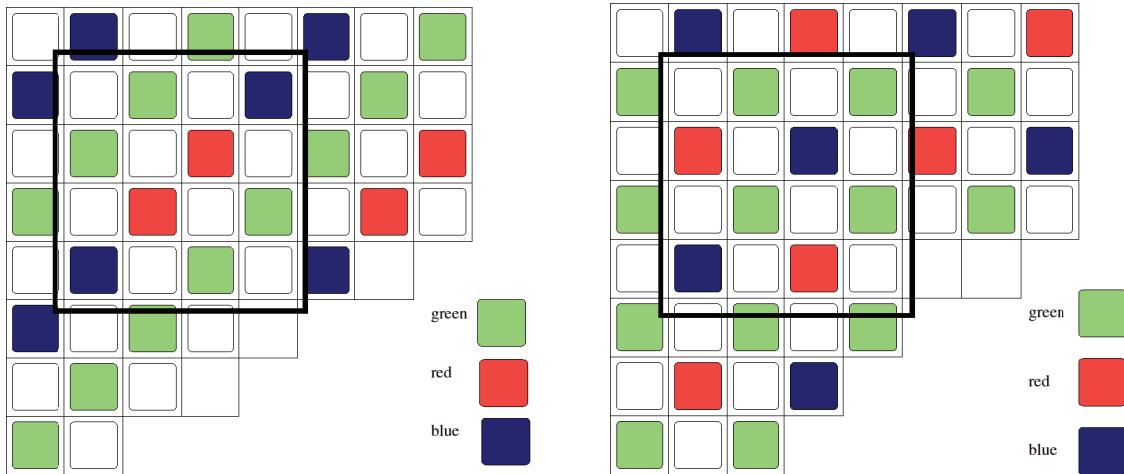


Figure 5.20. Two configurations offered by Kodak to replace the Bayer filter array: sites in white have a panchromatic sensitivity supposed to give a better smoothness to the image

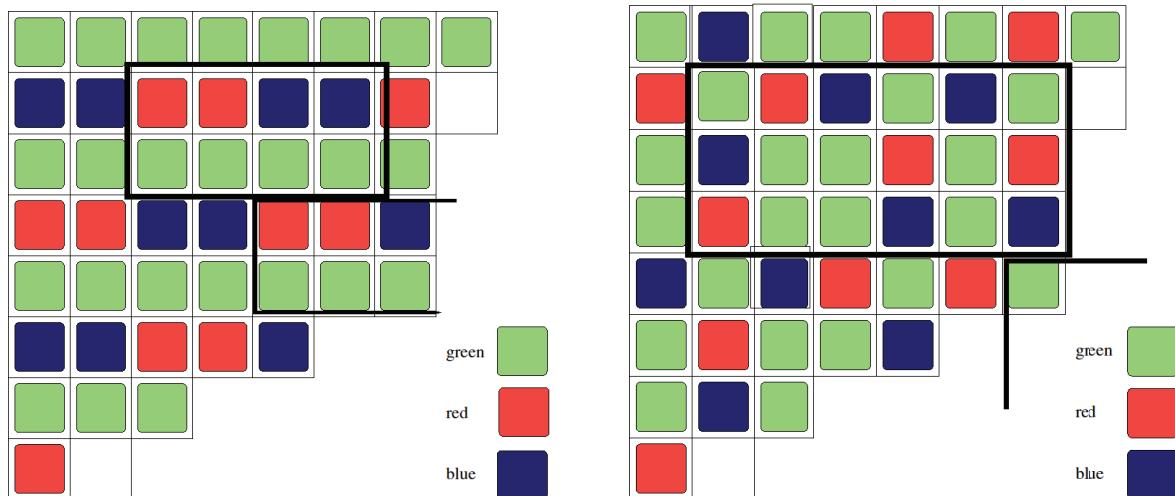


Figure 5.21. Configurations proposed by Fujifilm using original filter arrays: left, the ESR configuration uses a basic 4×2 pattern (in a quincunx) originating from a reconfigurable matrix (in the case, the array is rotated by 45° with regard to this representation, the neighboring pixels of the same color are coupled 2×2 to provide signals with stronger dynamics or a better signal-noise ratio). On the right, the X-Trans array (6×3 basic pattern on a hexagonal mesh) which made its appearance in 2013 and whose first objective is to overcome the moiré phenomena on high-frequency periodic textures



Figure 6.3. Three different noise levels affecting an image (detail) taken with high sensitivity (6,000 ISO). Left: no treatment applied to the sensor output. The noise is highly colored, due to the Bayer matrix which measures the signal. Center: noise reduction applied during demosaicing, followed by Bayesian filtering using local windows. Right: the same treatment, but using a non-local Bayesian filter

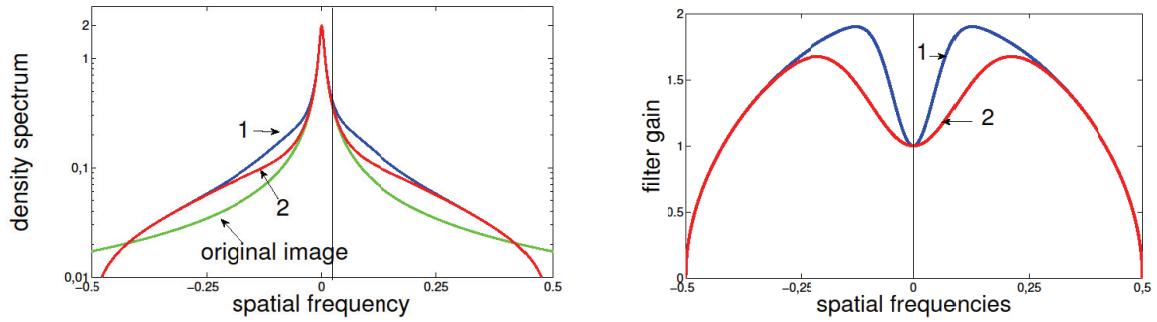


Figure 6.11. Unsharp masking: power density spectrum of an image (logarithmic scale), before and after filtering, using two different types of unsharp masking filters. Right: transfer functions of the two filters, in the form $[1 - \alpha \exp(-au^2)]$, multiplied by a function $\sqrt{\cos(\pi u)}$, which cancels out at the limits of the bandwidth in order to remove aliasing effects

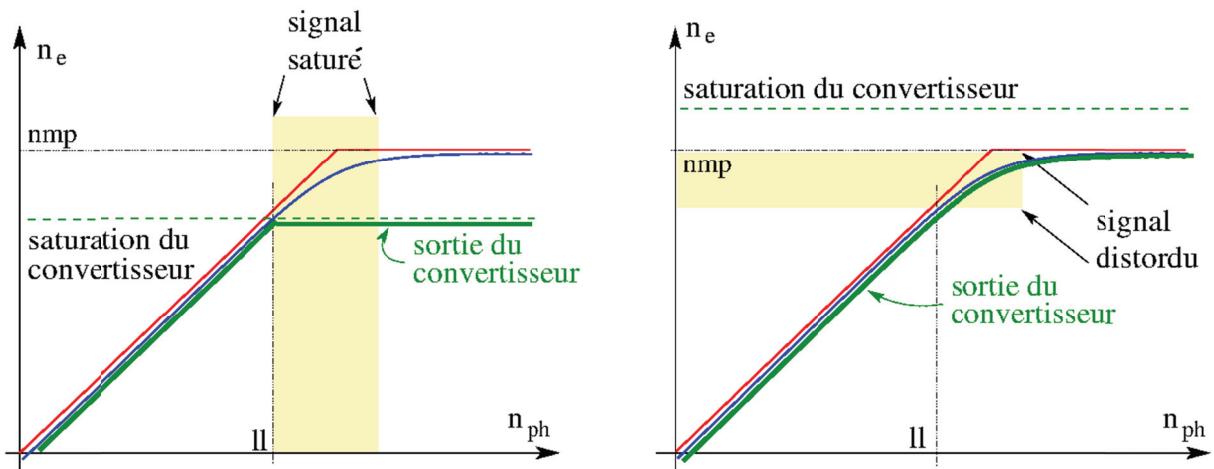


Figure 7.3. Sensor saturation: number of photons (n_{ph}) converted into electrons (n_e). In red: ideal photon-electron conversion; in blue, conversion in practice; in green, converter output. Two typical cases may be seen. Left: the converter is saturated first. Not all of the photons received by the sensor will be used to make up the image, but linearity will be respected in the converter dynamics (ll = limit of linearity). Right: the converter covers the whole capacity of the photosite (mcn = maximum carrier number), but now needs to take account of a zone of nonlinearity at the point of saturation

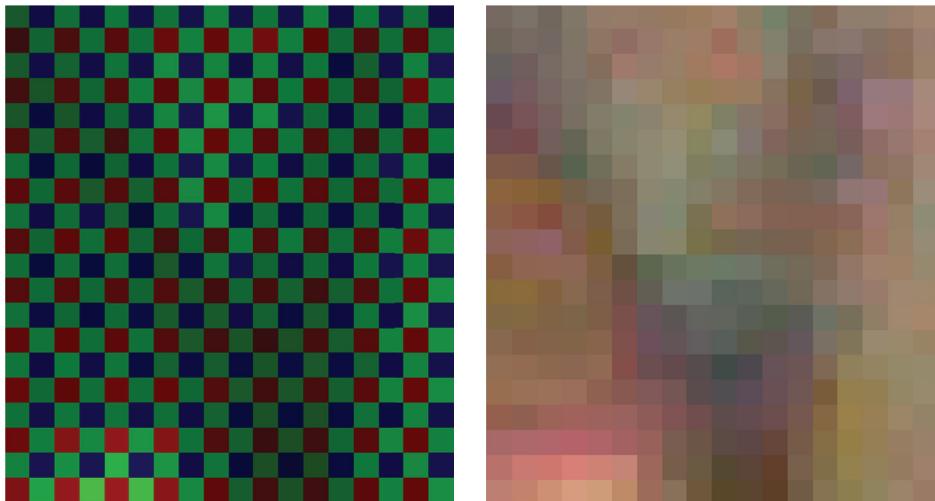


Figure 8.1. Left: a small section of the color image shown in Figure 8.2 (the top left corner), as recorded by the Bayer matrix and represented in a native (RAW) file, where the three channels, R, G and B, are stacked. For each pixel, only one of the R, G or B values is non-null. Right: the same image, reconstructed by demosaicing



Figure 8.2. Left: a color image. Center: the corresponding RAW file, represented in grayscale (the R, G and B pixels are interwoven, but have been reduced to gray levels, each coded on one byte). Right: the luminance image obtained from the color image. We see that the geometry of the two images on the left is identical, as the R, G and B channels are layered, but undersampled. The central image also shows the staggered structure characteristic of the Bayer mask, which is particularly visible on the face, which has low levels of green. This staggered structure is even more evident in Figure 8.1, which is highly magnified. On the lips, we see that the absence of green (which is twice as common in the samples than red or blue) results in the presence of a dark area, while the greenish band in the hair appears more luminous than the hair itself, which is made up of red and blue. The luminance image does not show this structure, as the RGB channels have been interpolated



Figure 8.5. Magnification of an area of an image encoded using JPEG. The top left image is the non-encoded original. The whole image has then been compressed by factors of 5.4 (top right), 12.1 and 20 (central line), and 30 and 40 (bottom line)



Figure 9.10. Blurring is an important aspect of the esthetic properties of a photograph. Left: spatial discrimination of objects exclusively due to focus. The choice of a suitable diaphragm is essential in softening contours. Right: the lights in the background are modulated using the optical aperture, giving a classic bokeh effect

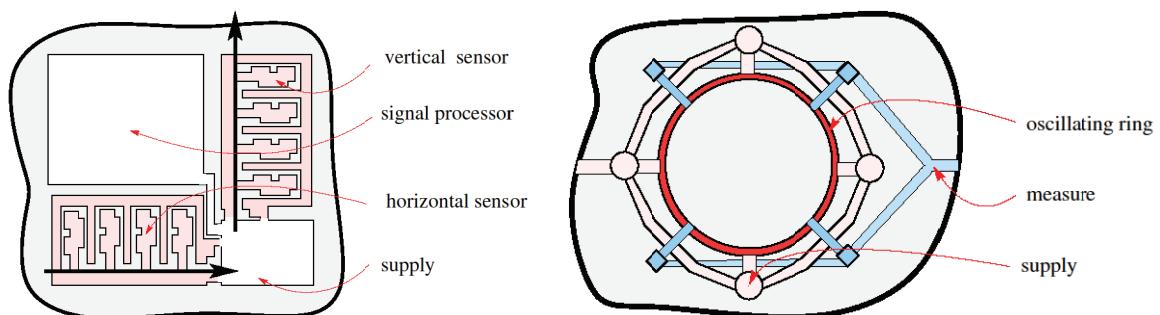


Figure 9.11. Left: two linear accelerometers using MEMS technology, made from interlinking combs, integrated into the same housing. These devices are used to determine translation movement in the plane of the host circuit. Right: a gyroscope, consisting of a vibrating ring (red, center). A circuit (shown in pink) is used to maintain vibration of the ring, while another circuit (in blue) measures deformation of the ring due to rotational movement. These two types of captors are produced from raw silicon



Figure 10.1. Denoising an image. Left: detail of an original image, with noise. Center: image after denoising using a bilateral filter. Right: image after denoising using non-local Bayesian filtering

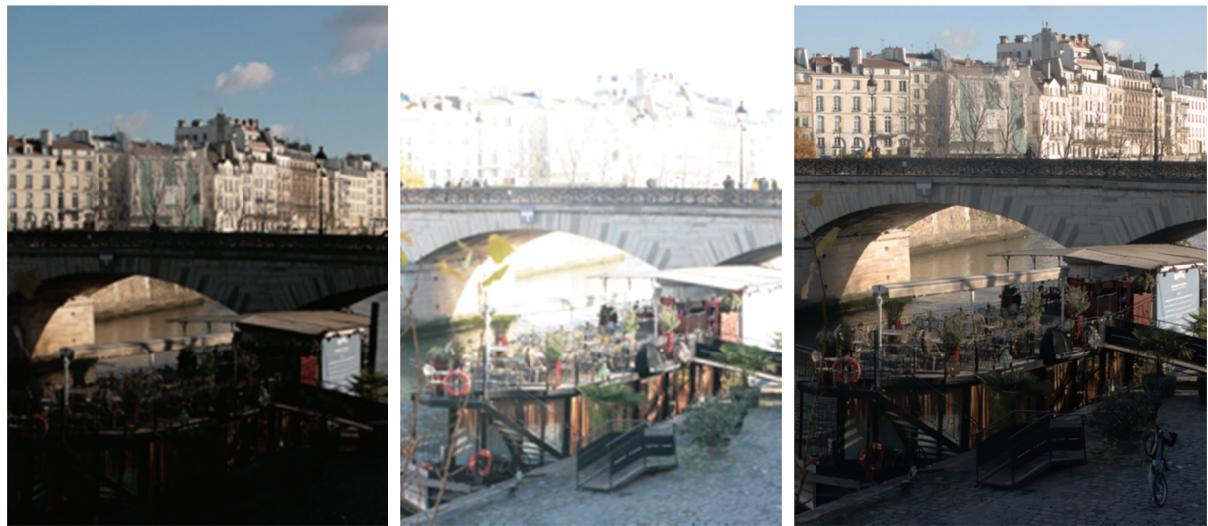


Figure 10.10. HDR treatment of four images (only the two extreme images are presented, with the least exposed on the left and the most saturated in the center). The result of this treatment, for the shared zone of the images, accounting for movement and compressed onto 8 bits, is shown on the right. Taken from [AGU 14a]



Figure 10.15. Flutter-shutter: example of image processing, taken from [TEN 13]. Left: the image recorded using a flutter-shutter, with a total movement of 52 pixels. Right: the corrected image. This second image is very close to the original, with the exception of the small vertical ripples in the sky on the left-hand side