Digital Camera Processing
By Dick Maybach, Member, Brookdale Computer Users Group, NJ
June issue, BCUG Bytes
www.bcug.com
n2nd (at) att.net

In the March and April 2018 articles (available at http://www.bcug.com/index.html), we looked at digital camera anatomy and the processes it uses to capture an image. The camera then applies significant processing to correct defects in the image and enhance it.

**Distortion**

Barrel distortion, shown in Figure 1, is a lens effect which causes images to be "inflated". It typically occurs at the wide-angle setting of a zoom lens and is most visible in images with perfectly straight lines, especially when they are close to the edge of the image frame.

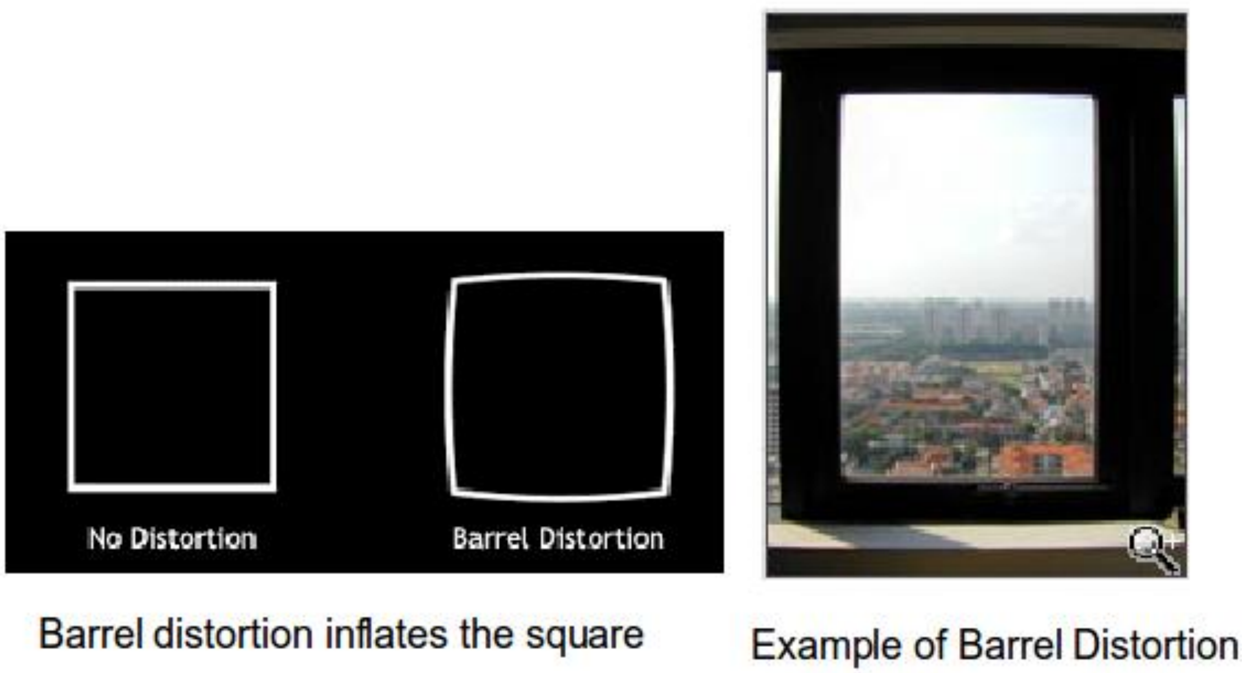


Figure 1. Barrel Distortion.

Pincushion distortion, shown in Figure 2, is the opposite of barrel distortion and causes images to be pinched at their center. It too is most commonly associated with telephoto zoom lenses, typically at their high-magnification end. Like barrel distortion, it is most visible in images with perfectly straight lines, especially when they are close to the edge of the image frame.

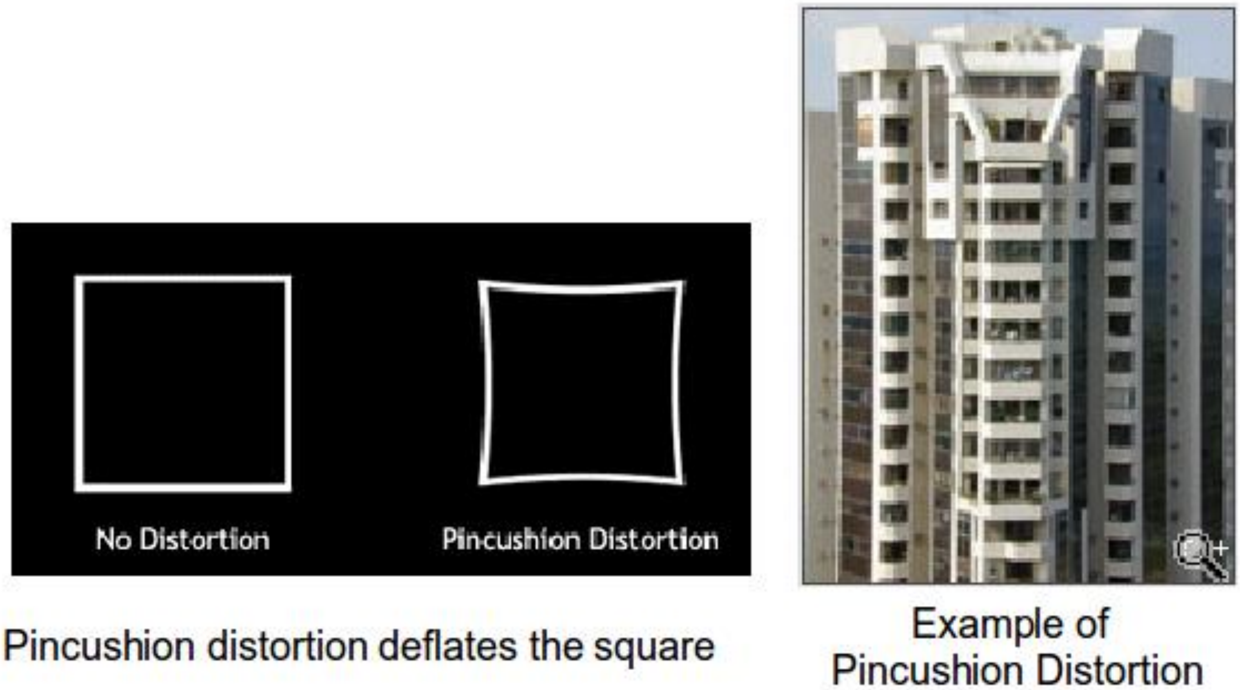


Figure 2. Pincushion Distortion.

In the film-camera era, distortion was corrected with additional lens elements, making them large, heavy, and expensive. Fortunately, if you know the lens settings, you know what distortion will occur, and this is how digital cameras correct it. They capture the image with distortion but move the pixels around to correct it before they record. Distortions of these types are functions of the lens, and interchangeable lenses have their own processors, which transmit the distortion characteristics of each image to the camera’s main processor.

Another lens problem is chromatic aberration or "color fringing" caused by the camera lens not focusing different wavelengths of light onto the exact same focal plane (the focal length varies with wavelength) and/or by the lens magnifying different wavelengths differently; see Figure 3. Chromatic aberration is visible as color fringing around high-contrast edges and occurs more frequently around the edges of the image frame in wide angle shots. The amount of this can be predicted from the lens settings, so correcting it is possible, although not always done.

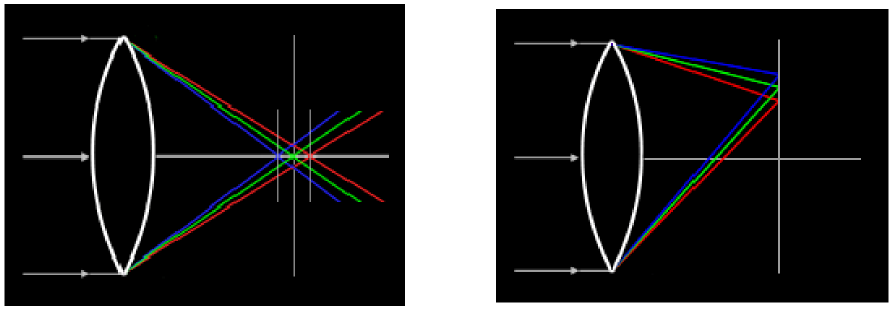

Longitudinal or Axial Chromatic Aberration -- Focal length varies with color wavelength

Lateral or Transverse Chromatic Aberration -- Position varies with color wavelength

Figure 3. Chromatic Aberration.

**Bayer Color Filter Array**

Each pixel on a digital camera sensor contains a light-sensitive photo-diode, which measures the brightness of light. Photo-diodes are monochrome devices, unable to sense color. Therefore, a mosaic pattern of color filters is positioned on top of the sensor to allow only red, green, or blue light to illuminate each pixel. The most common filter used in digital cameras is the GRGB Bayer Pattern, named after a Kodak engineer. The result is called a color filter array, shown in Figure 4. By breaking up the sensor into red, blue and green pixels, it is possible to get enough information in the vicinity of each pixel to make an accurate estimate of the true color there.

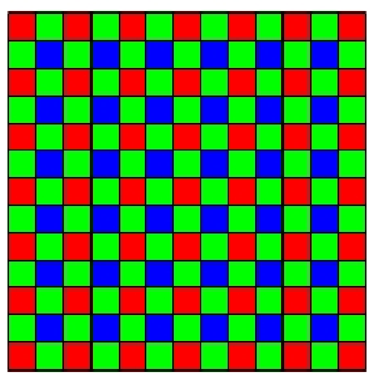

Figure 4. Bayer Filter.

In the Bayer filter pattern, the colors are not evenly divided; there are as many green pixels as there are blue and red combined, because our eyes are more sensitive to green detail than detail of the other colors. The advantages of this method are that only one sensor is required, and all the color information (red, green and blue) is recorded at the same moment. The raw output from a sensor with a Bayer filter is a mosaic of red, green and blue pixels of different intensities. After a raw image has been obtained from a photo-sensor with a Bayer pattern of color filters, it must be converted into standard format (usually sRGB or Adobe RGB), where each pixel is a mixture of red, green, and blue. This process is called demosaicing.

**Demosaicing**
Bayer demosaicing is the process of translating a Bayer array of primary colors into a final image which contains full color information at each pixel. One way of making this conversion would be to group four Bayer pixels onto a single full-color one. This would, of course, reduce, the horizontal and vertical resolutions, each by a factor of two. Effectively what we have done is to define a set of 2x2 boxes, use the data from all four pixels to calculate a single full-color pixel, and place it in the center of the box. These are shown as black dots in Figure 5.

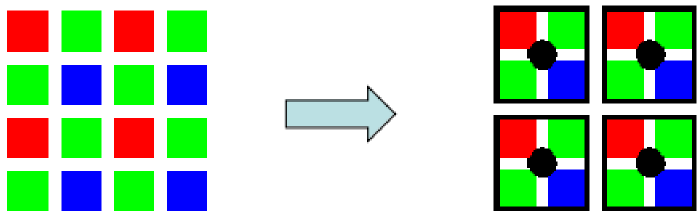

Figure 5. A Simple Demosaicing Process.

We improve the resolution by shifting our boxes one pixel to the right and computing a new set of full-color pixels also shifted one pixel to the right. Next, we could shift our boxes one row down and repeat. The result would be a set of the same number of full-color pixels as the total of all those in the Bayer array. Practical demosaicing is more complex and uses more than four Bayer pixels to calculate each full-color one. This doesn’t work well near the edges, since there are no longer pixels on all sides, but this is easily solved by discarding all the edge pixels. Because each full-color pixel is calculated from a group, there is some blurring of the image, which must be corrected by sharpening later.

**White Balance**
Pure white light is an equal mix of waves from all parts of this visible spectrum. In reality, light is seldom such an even mix. The color that illuminates the objects we photograph varies with the time of the day and with the light source e.g., candlelight, electric bulbs, and fluorescent lighting. Sometimes there are a greater number of the longer waves in the light, causing images to appear reddish. At other times there are more short waves, producing bluish light. This means that the colors, we see as reflections, also contain varying qualities of color. In Figure 6 the wall behind the lamp is white. In the image on the left, the camera has estimated the color of the white (called the white balance). The result is wrong, because the scene was lit by an incandescent lamp, which has a powerful orange color cast. In the photograph to the right, the white balance program was set to Tungsten, which fits the light given by electric bulbs and is more correct than the camera’s automatic estimate.

Figure 6. White Balance.

The tinge a "white" light has is described as its color temperature and is measured in degrees Kelvin (K). Figure 7 shows some color temperatures for common light sources: It might seem counter-intuitive that bluish light has a higher temperature than reddish light (since culturally we associate red with heat and blue with cold), but bluer light does contain more energy. The Kelvin scale was determined by progressively heating a black object so that it glowed red, then white, then blue. The temperature of a light source will affect the appearance of a scene or of colored objects. Our eyes and brain tend to compensate for this, but when, for example, a more objective device like a camera captures a candle-lit scene, the resulting photographs will often appear much too orange.

| Light Sources and Their Approximate Kelvin Values | | |
|---|---|---|
| Candle | 2000K | |
| Sunrise or Sunset | 2500K | |
| Standard household light bulbs | 3000K | |
| Noon on a sunny day "Daylight" | 5500K | |
| Electronic flash | 6000K | |
| Overcast sky | 7500K | |
| Blue sky | 12000K | |

Figure 7. Typical Color Temperatures.

**Compression**

The most commonly used digital image format is JPEG (Joint Photographic Experts Group). Universally compatible with browsers, viewers, and image editing software, it allows photographic images to be compressed by a factor 10 to 20 (compared to the uncompressed original) with very little visible loss in image quality. Figure 8 shows the effect of varying JPEG quality.

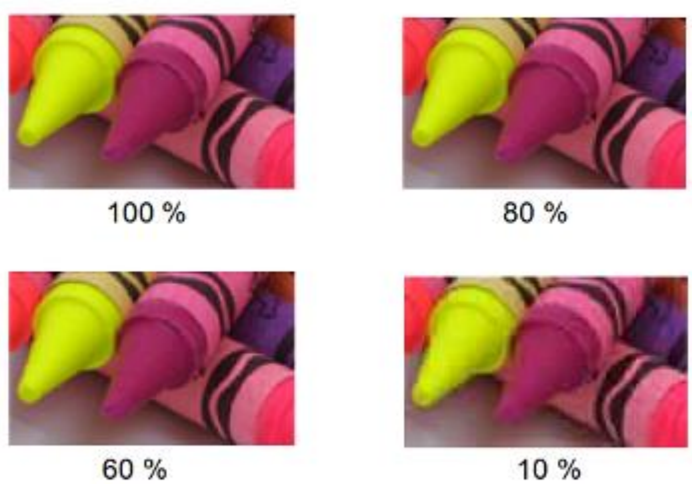


Figure 8. Effect of JPEG Compression.

100% quality JPEG images are very hard to distinguish from the uncompressed originals, which would typically take up 6 times more storage space. 80% quality JPEG looks still very good, especially when bearing in mind that this image is 2 times enlarged and that the file size is typically 10 times smaller than the uncompressed original. Notice some deterioration along the edges of the yellow crayon. Most digital cameras will use a higher quality level than 80% as their highest quality JPEG setting. If you look carefully at the 60% image, you will notice some JPEG squares and "hair" artifacts around the edges. It is a good trade-off because the file size is typically 20 times smaller than the uncompressed original. The 10% image shows serious image degradation. The only benefit of this low-quality level is that it illustrates what JPEG is doing in a more subtle way at higher quality levels. It is unlikely you will ever compress this aggressively.

**Correction and Enhancement**

Figure 9 shows an image captured by the sensor before any processing has been applied. Because half the pixels in a Bayer array are equipped with green filters, the image is quite green. It's also dark, because digital cameras deliberately underexpose to avoid saturating the pixels that record highlights. The next few screens will show the effects of processing.

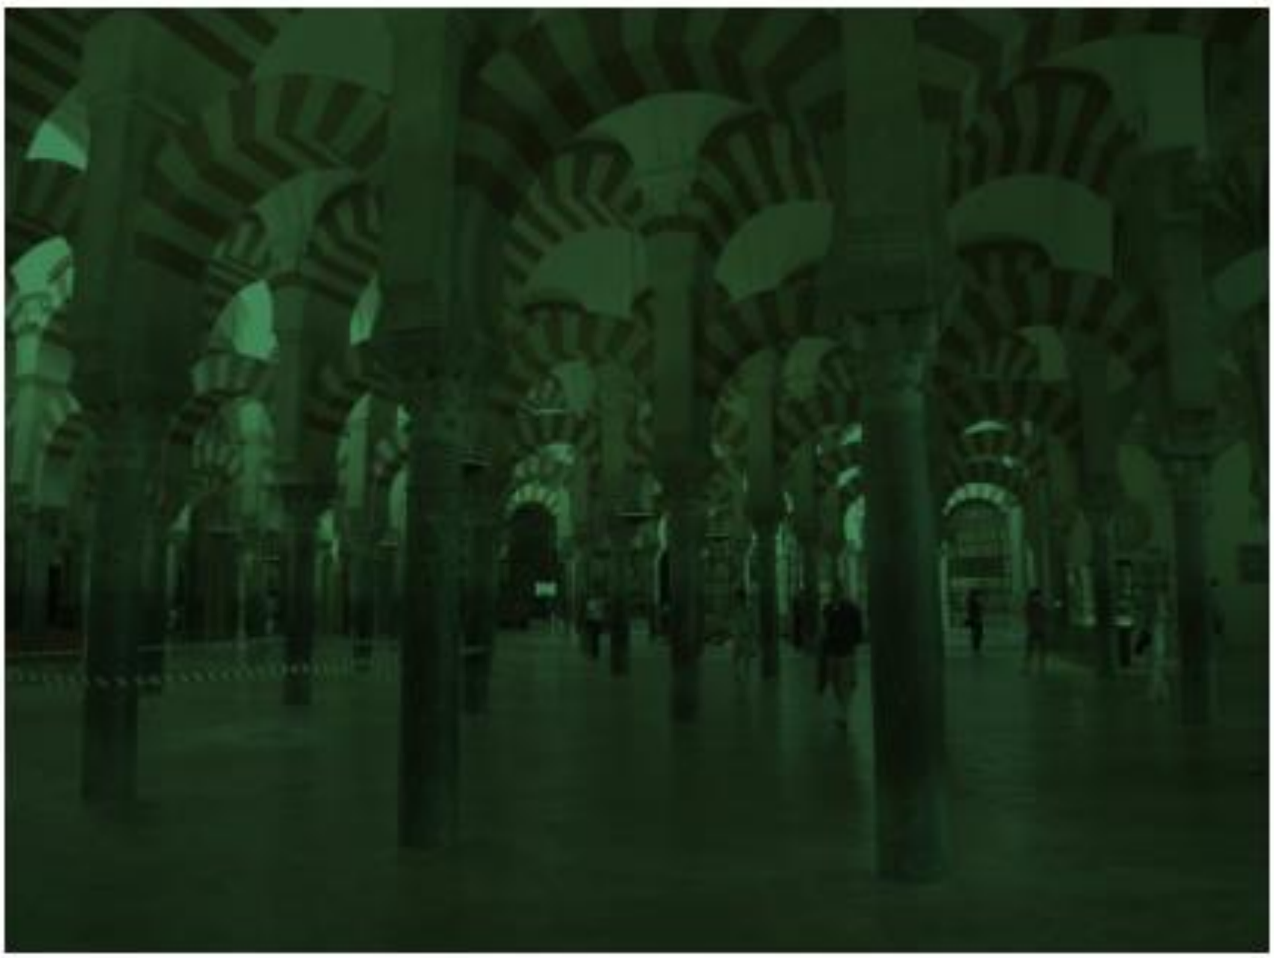

Figure 9. Raw Image as Captured by a Sensor.

Demosaicing converts the clusters of green, red, and blue pixels, as shown on the left of Figure 10. (The right shows the raw image before this processing.) The result is a set with each pixel having a combination of all three colors. The picture is still quite dark, as we haven’t yet compensated for the underexposure.

Figure 10. Image Before (Right) and After (Left) Demosiacing.

On the left of Figure 11 the exposure has been corrected. However, there is now too much red, because this picture was taken indoors, where the light is far less blue than outdoors.

Figure 11. Image Before (Right) and After (Left) Exposure Correction.

Your camera analyzes each picture and guesses what ambient light was present, or you can tell it. In Figure 12, I've used the camera's estimate, and the colors are truer.

Figure 12. Image Before (Right) and After (Left) White Balance Correction.

Lenses, especially zooms, have distortion. In this case, there was some barrel distortion, which has been corrected in the top half of Figure 13. Note the pinching of the frame edges and the disconnects within the pillars This can be corrected optically by adding more lens elements, but this is costly. It's more cost effective to correct the distortion with software. This requires a database of the corrections needed for each focal length. Cameras that use interchangeable lenses store this data in the lens processor. Other corrections are possible, for example, for noise, chromatic aberration, and blur, but these are difficult to see, and I'll skip over them.

Figure 13. Image Before (Bottom) and After (Top) Distortion Correction.

The result is somewhat dull. Most cameras will correct this by increasing the contrast and color saturation, as shown on the left in Figure 14. There is a trade-off between realism and drama, and many cameras have a setting to adjust this effect. When you think of all your camera must do to capture and store a realistic image, it's no wonder that it has far more processing power than desktop computers of just a few years ago. This is true of even the simplest point-and-shoot.

Figure 14. Contrast and Saturation Correction (Left).

The processing to develop a single image is complex and occurs repeatedly when recording videos and viewing scenes before you press the button, the computing power residing in a digital camera, even a relatively simple one in a cell phone, is quite remarkable. Hopefully, this series has given you some appreciation of it.