

*INTRODUCTION TO
RAW FILES OR
DIGITAL NEGATIVES*

1



CAMERA: CANON 100

Our goal in this book is to demonstrate a reliable workflow that will help you achieve optimal image quality from your digital camera using the RAW file format. After capturing RAW images, you can use your computer (rather than your camera) to convert them to JPEG or TIFF.

To help you better understand what you'll be doing when you work with the RAW converter software covered in later chapters, this chapter will explore how the camera creates a JPEG file. In the process, we'll cover some of the limitations of in-camera file conversion, and we'll outline a basic strategy for beginning your workflow with the best possible source image.

1.1: How the Camera Creates JPEGs

JPEG (Joint Photographic Experts Group) is the image format most commonly used by today's digital cameras. It provides reasonably good image quality, but as you'll see, it has some limitations. The obvious one is that its compression format, although excellent, is lossy; some information is always thrown away when an image is compressed. (There are lossless JPEG formats, but they are rarely used in cameras.) And, even when the JPEG compression is low, the image still degrades slightly.

A more significant problem is that before a camera converts an image to JPEG format, the image must undergo extensive processing within the camera. This processing includes color and exposure correction, noise reduction, and sharpening. Because these adjustments are made in the camera, your ability to make further post-processing corrections is limited.

The upshot is that JPEG compression works best for images that don't need further substantial post-processing, or when such post-processing is prohibited. However, if you demand high quality from your work, you will find that it is the rare image that does not need some post-processing. 📷



NOTE:

Many photographers try to get the best possible quality from their cameras by saving their images as TIFF or as a vendor-specific RAW file format (if available). As you will see, RAW file formats allow the greatest post-processing flexibility.

HOW IN-CAMERA CONVERSION WORKS

All new digital cameras capture color photos, right? Well, not exactly. While you ultimately get color prints from a digital camera, most modern digital cameras use sensors that record only grayscale (brightness or luminance) values. (The Foveon x3 sensor, digital scanning backs, and multishot digital backs are exceptions.)

For example, say you want to photograph a box of Crayola crayons (Figure 1-1).

A grayscale sensor would see the picture as shown in Figure 1-2; that is, it would see only shades of gray.

But how do you use a grayscale sensor to capture color photos? Engineers at Kodak came



FIGURE 1-1: Full color sample target




FIGURE 1-2: Grayscale picture seen by the sensor




FIGURE 1-3: Bayer Pattern achieved by a matrix of color filters

up with the Color Filter Array configuration illustrated in Figure 1-3. This configuration is called the Bayer Pattern after the scientist who invented it back in the 1980s. (Other pattern variations are used, but this is the basic technology used in most CCD and CMOS sensors.)

The yellow squares in the grid shown in Figure 1-3 are the photoreceptors that make up the sensor; each receptor represents one pixel in the final image. Each receptor sees only the part of the light that passes through the colored filter just above the sensor element (either red, green, or blue).

Notice that 50 percent of the filter elements (and thus the receptor elements) are green and only 25 percent each are red and blue. This pattern works because the human eye can differentiate many more shades of green than it can red or blue ones (which should be no surprise when you consider the number of shades of green in nature). Green also covers the widest part of the visible light spectrum.

Each receptor in the sensor captures the brightness value of the light passing through its color filter (Figure 1-4).

Each pixel, in turn, contains the information for one color (like a tile in a mosaic). However, we want our photo to have full color information (the R, G, and B channels) for every pixel. How does that magic happen? Here's where a software trick comes into play: a process called demosaicing or color interpolation adds the missing RGB information by estimating information from neighboring pixels.




FIGURE 1-4: Color mosaic seen through the color filter

Demosaicing and Color Interpolation

Demosaicing, then, is the method of turning the RAW data into a full color image. A good demosaicing algorithm is quite complicated, and there are many proprietary solutions on the market.

The challenge is to resolve detail while at the same time maintaining correct color. For example, think of capturing an image of a small, black-and-white checker pattern that is small enough to just overlay the sensor cells, as in Figure 1-5.

White light consists of red, green, and blue, and the white squares in our example pattern correspond exactly to the red- and blue-filtered photoreceptors in the sensor array. The black squares, which have no color information,

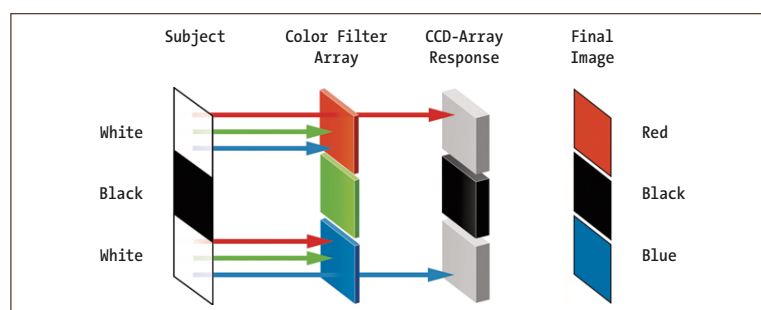


FIGURE 1-5: CCD/color-mosaic sensor with color interpolation errors




FIGURE 1-6: Foveon sensor with no color interpretation problems



NOTE:

These test photos show an extreme situation because a Bayer sensor cannot really capture the transition from blue to red at a pixel level. While the failure is less dramatic in real-world photographs, it is still visible and definitely not something to be ignored.

correspond to green-filtered photoreceptors. So for the white squares that align with red photoreceptors, only red light passes through the filter to be recorded as a pixel; the same is true for the blue photoreceptors.

Color interpolation cannot correct these pixels because their neighboring green-filtered photoreceptors do not add any new information, so the interpolation algorithm would not know whether what appears to be a red pixel really is some kind of “red” (if the white hits a red filter) or “blue” (if the white hits the blue filter).

Contrast this with the Foveon sensor technology illustrated in Figure 1-6. Instead of the Bayer Pattern where individual photoreceptors are filtered to record a single color each, the Foveon technology uses layers of receptors, so that all three color channels are captured at the

same photosite. This allows the Foveon sensor to capture white and black correctly without the need for interpolation.

The resolution captured by the Bayer sensors would decrease if the subject consisted only of red and blue shades, because the pixels for the green channel could not add any information. In the case of monochromatic red or blue colors (those with very narrow wavelengths), the green sites get absolutely no information. But such colors are rare in real life and, in reality, even when the sensor samples very bright and saturated red colors, information is recorded in both the green and (to a much lesser extent) blue channels.

The problem in our example above is that in order to correctly estimate the color, we need a certain amount of spatial information. If only a single photosite samples the red information, there will be no way to reconstruct the correct color for that particular photosite.

Figure 1-7 illustrates a test we did in a studio to demonstrate the loss of resolution with a Bayer sensor when capturing monochrome colors. Notice how blurry the text in the Canon image is compared to that of the Sigma image. [⚡]

Some of the challenges that face interpolation algorithms include image artifacts, like moirés and color aliasing (shown as unrelated green, red, and blue pixels or resulting in discoloration). Most cameras fight the aliasing problem by putting an antialiasing (AA) filter in front of the sensor (which actually blurs the image and distributes color information to the neighboring photosites). Of course, blurring and high-quality photography don’t usually go together, and finding the right balance between blurring and aliasing is a true challenge for camera design engineers. (In our experience, the Canon 1Ds does this job well.)

After antialiasing, an image needs to have stronger sharpening applied in order to re-create much of the original sharpness. (To some extent, AA-filtering degrades the effective resolution of a sensor; therefore, some strong sharpening is typically needed later, during the RAW workflow.)

The mission of creating a high-quality image from the data recorded by a sensor is a complicated one, but it works surprisingly well. Every technology has to struggle with its inherent




FIGURE 1.7: Fooling a Bayer sensor

limitations and digital photography can beat film today in many ways because film has its own limitations. 🌸

THE LIMITATIONS OF IN-CAMERA PROCESSING

For any given digital camera, the RAW data is all the data for grayscale brightness values captured on a chip. To produce a final image, this RAW data needs to be processed (including demosaicing) by a RAW converter. In order to produce JPEG images, the camera must have a full RAW converter embedded within its firmware.

As you've already seen, one effect of in-camera conversion to JPEG is artifacts caused by lossy compression, but camera-produced JPEGs have other limitations as well:

- Although most sensors capture 10- to 14-bit color (grayscale) information, only 8 bits are used in the final file. (JPEG can't encode more than 8 bits per color channel.)
- The in-camera RAW converter can use only the camera's own, limited computing resources. Good RAW conversion can be very complex and computer intensive. It is much more efficient to have the host computer convert the image than it is to use the onboard ASIC chip commonly used today. Additionally, the in-camera chip can't be upgraded, so as software technology evolves, the difference in efficiency will only increase.
- White balance (WB), color processing, tonal corrections, and in-camera sharpening are all applied to the photo within the camera. This fact limits your post-processing capabilities because a previously corrected image must be corrected again, and the more a photo is processed (especially 8-bit), the more it can degrade.

RAW FILE FORMATS

The advantage to working with RAW file formats directly is that they essentially store only RAW data (although they also contain an EXIF section, which holds additional metadata that describes properties such as camera type, lens used, shutter speed, f-stop, and more). Fortunately, you can perform all the processing that would be

done in the camera to convert a file to JPEG or TIFF (including white balancing, color processing, tonal/exposure correction, sharpening, and noise processing) on a more powerful computing platform. This offers several advantages:

- No JPEG compression. You can work directly with the RAW data.
- You can take full advantage of the 12-bit color information (10 to 14 bit). This becomes particularly significant if you need to make major corrections to the white balance, exposure, or color. When processing an image, you lose bits of image data because of data clipping (which accumulates over multiple steps). The more bits you begin with, the more data you'll have in your final corrected image.
- You can use very sophisticated RAW file converters such as Adobe Camera Raw, Pixmantec's RawShooter Essential, Apple's Aperture or Phase One's Capture One DSLR. Chapters 4 through 7 provide hands-on guidance for using these tools.
- You can do color correction (white balance).

Working with JPEGs created in-camera is like working with images produced by a Polaroid camera (where you simply shoot and receive your processed image immediately). Working with RAW files is more like working with a traditional film negative that can be developed and enhanced in the darkroom. RAW converters mimic the film development process, and because you can always return to your original RAW file and process it again, you aren't limited by the technology built into today's cameras or even today's software. Over time, improved RAW file converters will produce even better results from the same data.

All in all, shooting in RAW gives you much greater control while processing your images.

The TIFF Option

What about setting your camera to save images as TIFF files? Saving as TIFF files solves only the lossy compression issue because the images are still converted to 8-bit inside the camera. Too, most TIFF files are larger than RAW files (RAW files hold



NOTE:

Keep in mind that lost detail cannot be recaptured regardless of the fancy sharpening methods used.

Sharpening does, however, do quite a good job, and digital cameras work fairly well.

only one 12-bit grayscale value per pixel), and they don't offer the flexibility and control benefits that RAW does. An 8-bit in-camera processed TIFF file is only slightly better than a high-quality/high-resolution JPEG.



1.2: The Digital Negative/Slide



The files created on your computer when shooting RAW are often called digital negatives. You should keep these RAW (or original JPEG or TIFF) files even after converting them because they hold all the information captured when you took your original shot. You might want to revisit them when:

- You've improved your own digital workflow (which is very likely over time).
- Better RAW converter software becomes available. We have seen many improvements over the last four years and expect more to come.
- You lose your derived files.

A RAW file is like a latent image, and RAW converter software can bring detail out of overexposed or underexposed shots. The one big difference between film and digital photography is that you can perform multiple kinds of development on your digital images over time, and you can return to them over and over again.



1.3: Some Strategic Reasoning



Just as many paths lead to Rome, there are many ways to shape your digital workflow. The most appropriate way for you will depend on the kind of photographs you shoot, how you intend to use or reproduce your images, your equipment, and your personal preferences. To set up a workflow that suits your work best, you may need to try different variations before you finally settle on one (which you may still adjust for special cases).

Several of the steps we follow when processing an image (including white balancing, sharp-

ening, enhancing contrast, noise reduction, and enhancing saturation) can be performed at any one of three different stages:

- Inside the camera (some of the operations mentioned above, even if you shoot RAW files)
- Inside your RAW converter
- When using Photoshop or a Photoshop plug-in

Performing an operation at an early stage doesn't necessarily mean that you cannot perform it again at a later one. For example, there will be times when it will make sense to perform some sharpening globally with the RAW converter and then add more sharpening when you tweak the image with Photoshop, which allows you to sharpen either the whole image or certain edges or areas.

You may even want to perform some final sharpening when you are preparing a dedicated form of output. (For example, an image printed using an inkjet printer or offset printing will require more sharpening than one that is to be presented on-screen or printed with a lightjet printer on photographic paper by a photo service.) In-camera RAW converters and software RAW converters allow only global changes to the whole image, while Photoshop allows you to make selective corrections using selections, masks, layers, and filters.

The general rule is that corrections done in the camera or (with more control) in the RAW converter, will do less to degrade the quality of your images than corrections made in Photoshop. This is especially true when correcting exposure, which should be done within the camera; it is much better to correct exposure values with the RAW converter than to do so within Photoshop.

WHITE BALANCE

When you shoot RAW files, make sure to set the proper white balance in the camera, because this value will be used as the default starting point for your white balance inside the RAW converter. However, this is only a starting point and may be changed without losing color quality. White balancing is one of the most important tasks in a

RAW converter and should usually be done there (as you'll see demonstrated in the chapters on individual RAW converters). [⚡]

SHARPENING, SATURATION, CONTRAST ENHANCEMENTS

Apart from the exposure and white balance settings, you should deactivate all other settings (such as sharpening, saturation, and contrast enhancements in the camera) or set them to the lowest possible value when you shoot RAW files. This is even more important if you shoot JPEG or TIFF. The only exception would be if you shoot JPEG or TIFF and don't want to do any post-processing in Photoshop.

UP-SAMPLING OR DOWN-SAMPLING

When scaling, we recommend using the camera's highest resolution and then performing up-sampling or down-sampling either in the RAW converter or in Photoshop. Photoshop and Capture One both support reasonably good up-sampling, but they offer only certain fixed sizes. To prepare an image for large-scale printing, you may find it helpful to do a rough up-sample in your RAW converter and a final sizing in Photoshop.

CROPPING

Adobe Camera Raw 3 (or newer), Nikon Capture, and Capture One all offer cropping. Cropping results in smaller image files and, as a result, faster processing.

1.4: Choosing a RAW Converter

Several RAW converters are available, and more are being released. Almost all DSLRs come with a native RAW converter, which might be a Photoshop plug-in, a stand-alone RAW converter, or both. For example, Canon cameras that offer the RAW format ship with both the Canon EOS Viewer Utility (EVU) and Canon Digital Photo Professional (DPP). If you use Photoshop CS1 (alias Photoshop 8) or CS2 (alias Photoshop 9),

or even Photoshop Elements 3 or higher, your software includes a good RAW converter (for Adobe Camera Raw, see Chapter 4). If you use a picture database or a good picture file browser (also referred to as a Digital Asset Management Systems or DAMS), it probably will have a RAW converter of its own. Stand-alone RAW converters include RawShooter by Pixmantec (covered in Chapter 5), Capture One by Phase One, Nikon's Capture, and Bibble (all covered in Chapter 7)—just to name a few.

The RAW converters that are part of a DAMS system (such as ThumbsPlus, iView Media Pro, and Extensis Portfolio) are primarily intended to produce a reasonable preview. Although they can be used to convert a RAW file to TIFF or JPEG, conversion is not their primary focus. Therefore, you should seriously consider converting your RAW file with one of the "real" RAW converters.

The range of formats and cameras supported by the various RAW converters differs. Nikon's tools, for example, support only Nikon cameras, and Canon's tools support only Canon's cameras. Adobe Camera Raw, Capture One, RawShooter, and Bibble support a wide range of formats, but check to be sure that they support your camera. Too, some converters are faster than others, and some offer better workflow integration. For this reason alone, you should read the descriptions of all the RAW converters covered in Chapters 3 through 7, download their trial versions from the Internet, test them, and settle on the one that best fits your needs and budget. You may also find it useful to use more than one converter, depending on the type of work and the images you have to process.



NOTE:

With cameras that support black-and-white shooting, you may consider using a black-and-white setting if you are sure that you don't want color. However, shooting color offers more control and more options. You can always convert color to black-and-white inside the RAW converter or in Photoshop (see Chapter 14).
