-
- SHARE

Exposing For RAW

There are special considerations to take into account if you're shooting RAW and you want to be sure that you're getting a proper exposure

By Andrew Rodney

You wouldn't think changing image capture from film to digital photography would require a new way to think about exposure, but it may, depending on how you use your digital camera. This is because a digital camera sensor behaves quite differently from how film and our human visual system respond to light intensity. Digital cameras record data in a linear fashion. The human visual system responds to stimuli in a nonlinear fashion. Example: You enter a pitch-black room and turn on a 100-watt light; you see a fixed amount of light intensity, as would a digital camera. You then turn on a second 100-watt light, thus doubling the light output. Because humans have a built-in nonlinear visual system, the room doesn't appear twice as bright. Because of the linear behavior of a digital camera's image sensor, how-ever, the camera would record the scene as twice as bright.

What Is RAW Linear Data?
The digital sensors in our cameras translate the intensity of light energy (photons) into an electrical charge, which is converted to a digital value. In Figure 1, you can see how a hypothetical camera sensor can record six stops of dynamic range, from shadow to highlight, using 12-bit encoding (0 to 4096 levels). The darkest tone the sensor can record at this low light level has a digital value of 0. The brightest and largest number of light energy this sensor can capture has a value of 4096.

In linear capture, the first stop of highlight data contains half of all the levels within the image, and the next stop, half the remaining number of levels and so on. The fewest number of levels recorded is in the last stop of shadow detail—only 64 of 4096! If the image is underexposed, even fewer levels are used to describe that last stop of tonal data. This is also where most camera noise is found, so the results of underexposure are more noise, less actual data. Even with no exposure whatsoever, you'll find noise generated from ambient radiation. Shoot a frame with your lens cap on and then examine the RAW data in your converter, perhaps pulling a steep curve or moving the exposure sliders to examine the data represented in the dark shadows. You won't see all black, level-zero pixels, but instead some random noise.

Depending on the dynamic range of the scene being photographed, there may be darker tones the sensor can't detect and brighter tones that exceed what the sensor can record due to exposure (4097+ photons). In a scene that has 6½ stops of dynamic range, it's possible that a half stop of brightness can't be recorded because this portion of scene brightness occurs beyond full-sensor saturation. All light values above full saturation can't be recorded unless the exposure is adjusted. Doing so would result in a half-stop less shadow data being recorded. You must attempt to fit the dynamic range of the scene within the dynamic range the capture device can record.

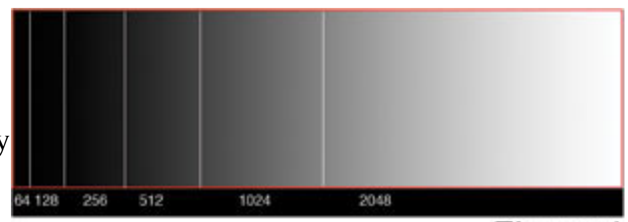

Figure 1

You can control the lighting ratio to fit this range, or expose to record either highlights or shadows, resulting in loss on the opposite end of the tone scale. The opposite extreme is actual overexposure, whereby the sensor can't record additional data in the highlights past the point of full-sensor saturation. Here, exposing for digital or film is the same: avoid overexposure that blows out highlight detail you hope to reproduce.

When you edit the RAW file in a converter, the software is operating on this linear-encoded data, which provides some useful advantages, such as powerful control over adjusting the tones in the first stop of highlight data. Some RAW converters can recover highlight data that would be difficult or impossible to alter in a nonlinear and gamma-corrected image. When film records a scene, it's necessary to follow the old practice “expose for the shadows, develop for the highlights.”

This brings us to the concept of Exposing To The Right (ETTR), that is, exposing for highlights and developing for the highlights. This ETTR recommendation came about on the Luminous Landscape Website several years ago, based on an interview with Thomas Knoll, one of the original authors of Adobe Photoshop and Adobe Camera Raw (www.luminous-landscape.com/tutorials/expose-right.shtml). Expose to place as much data within this linear-encoded RAW image without losing highlight values you wish to reproduce. Capture the most image data the sensor can record.

Is What You See What You Get?

ETTR presents a few problems, one being that the LCD camera preview, including the histogram and clipping indicators, isn't based on the linear RAW data. Instead, this preview is based on the rendered gamma-corrected JPEG your camera is set to produce, even if you don't save that JPEG and only shoot a RAW file! If your goal is to produce the best possible exposure for RAW, using the ETTR technique, the feedback on the LCD could steer you in the wrong direction. The camera uses its on-board processors to render color and tone from RAW to create a JPEG, which may be vastly different from what you hope to render from the RAW using a stand-alone converter. This is primarily why photographers shoot RAW. You want to control the color and tone rendering instead of accepting how the camera processes the RAW into a rendered JPEG. Yet you have no direct feedback about this RAW data when shooting.

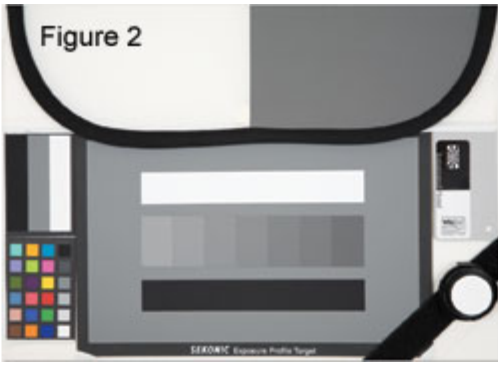

On a recent photo tour down the Amazon River with a group of photo students, I was rather surprised to see that several had set their exposure compensation for -1 stop because the color and tone rendering on the camera LCD "looked better." These students, all shooting RAW, were conducting the opposite of the ETTR technique, introducing more noise and less data, to produce a better-appearing LCD preview.

During this expedition, I decided to shoot a number of bracketed exposures and see if there was any correlation between proper exposure, based on ideal highlight data, and the appearance of the images on the LCD. Unfortunately, at least with my Canon EOS 5D, the "best" exposure often produced the worst-appearing LCD previews. Other than the possibility of truly clipping highlight data in RAW, this ugly LCD preview is one of the biggest disadvantages of utilizing the ETTR technique. You may have some camera controls that affect how the histogram and clipping indicators are plotted on your camera's LCD. Many cameras have "picture styles," which alter the JPEG rendering and therefore, the LCD previews. Testing is necessary to adjust these picture styles while viewing highlight clipping on your camera's LCD to more closely match the resulting RAW data clipping if present. Such settings may provide closer clipping previews, but don't expect a match to the real clipping of the RAW data. Plus, these picture styles will affect JPEGs you may shoot, but not the RAW data—not useful for those who wish to shoot RAW+JPEG.

When I set my 5D Picture Style from the default to -4 contrast, the clipping of highlights on the LCD didn't appear until I had "overexposed" 2⁄3 of a stop. This is better than the original default settings, yet in actuality, I was able to overexpose 1½ stops beyond what my meter suggested while fully retaining highlight data in my RAW file. The clipping indicators are still far from correctly describing what's happening to the RAW data.

Testing ISO And Exposure

Figure 2 shows a number of exposure and white-balance targets, evenly illuminated by two Balcar strobes within 1⁄10 of a stop in all four corners. A Sekonic L-758DR handheld flash meter was used to measure incident light for a base (normal) exposure. Images were shot first using this "normal" exposure for ISO 100, 400 and 800, then -1 stop and +2 stops overexposed in half-stop increments. I bracketed exposure by altering the strobe power rather than the *f*-stop, which introduced too many visual variables caused by depth of field and apparent lens sharpness. I was unable to test at ISO 1600 or underexpose at ISO 800 because of the power of the strobes. However, the three ISO settings do provide some useful information about ETTR on the RAW data.

Once I had a series of bracketed exposures, I examined each in my RAW converter of choice, Adobe Lightroom. The image exposed as recommended by the flash meter did appear properly exposed using the default Lightroom rendering settings. It was no surprise that the over-exposed images looked, well, over-exposed **(Figure 3)**.


Figure 3

I now needed to create a custom rendering setting in Lightroom for ETTR. Since images exposed to the right initially looked overexposed, I needed to alter the rendering controls to produce an acceptable color and tone appearance and create a new default preset that I called "normalize." Since Camera Raw and Lightroom utilize the identical processing pipeline, I can use this custom normalize default setting in either product. This is especially useful in Lightroom because I can select a custom preset upon importing my images **(Figure 3)**. I never see the ETTR images in Lightroom as overexposed. If you use another RAW converter, you should still be able to produce a custom setting that compensates for the initial "overexposed"-appearing rendering. Alter the necessary tone controls in the converter to produce a white with just a small amount of tone, then save that as a new default. The key is exposing such that white in the target is as close to, but not clipping, 255/255/255.

In Lightroom, two stops over the base exposure blew out the highlights, and no amount of exposure compensation could rectify this. The sensor was pushed to full saturation, a situation that should be avoided. The 1½-plus exposure bracket could be normalized by setting the Lightroom Exposure slider to -1.47. No other settings would be adjusted. I attempted to set the gray background of the Sekonic target to match the "normal" exposure while keeping the all-white targets below 100 percent pure clipping value.

For the ISO 400 image, I had to move the exposure slider in Lightroom to -1.68 to get the overall gray of the target background to read approximately the same as the normal version, while avoiding highlights clipping. I noticed that the overall contrast of the various images looked slightly different from each other. This indicates that boosting ISO may affect other areas of the tone curve.

For the ISO 800 image, I was able to use the normalize exposure preset used for ISO 400 to produce a close match to the base exposure. With this information, I was able to extrapolate the actual-sensor ISO for exposing images to place as much data in the first stop of highlight data without blowing out the highlights.

To better match all exposure versions visually, white balancing was necessary. I used the Lightroom WB tool on the neutral gray found on the WhiBal card in the same location for all images. The result produced only minor alterations throughout the image, while not all grays in all images produce the same numeric values. The various grays visually appear slightly different even though the values are within +/-2 to 3 percent as indicated by the Lightroom info readout.

When viewing the normalized “overexposed” images and the “normal” exposed images at 100 percent or higher, the ETTR technique always produced superior-quality images, but the degree was highly dependent on ISO. Progressively higher ISO set-tings resulted in better quality, yet in all cases, the differences of ETTR were visible. As the math suggests, less noise and often better shadow detail can be seen.

I still had to deal with the in-camera meter, which measures the reflected, not incident light, then come up with some compensation factor if I didn't want to resort to using an external incident meter, which is highly recommended. Reflective meters are calibrated to “see” or measure re-flectance, depending on make, of 12.5 percent or 18 percent gray. Aim such a meter at a white dog in snow and you end up with a gray dog. Point the meter at a black cat on a pile of coal and you end up with a gray cat.

Photographers still need to understand the limitations of a reflective meter based on what it's currently measuring and adjust accordingly. For example, an old analog film trick was to take a meter reading on your hand illuminated by the light you'd be shooting in, then open up one stop. Modern cameras have a number of metering modes that can measure multiple areas in an image and attempt to provide better exposures than a spot meter described above, yet the fundamental concepts are the same. Having produced a series of test exposures using an incident meter, I hoped to find the actual ISO sensitivity of the camera sensor based on the ETTR technique.

ETTR: Is It Worth It?
The practical implications of ETTR are to produce superior image data, especially with respect to noise in the shadows. The math, which suggests that ETTR produces better data, is undeniable, but does an ETTR technique have practical advantages?

For high-ISO shooting, it's clear from the test examples that noise is greatly reduced; however, shooting in this mode effectively provides a much lower initial ISO setting, so this technique might not be that useful. The noise-reduction advantages are less noticeable at lower ISO settings, and the disadvantage to this technique is the appearance of images displayed on the camera's LCD.

Getting acceptable-looking LCD previews and ideal exposure don't mix, not until the camera manufacturers radically alter the way they build the preview based on the RAW data being captured. The histogram provided is equally less useful because it's gamma encoded based on the camera-generated JPEG settings. Clipping indicators could be useful, assuming you test your camera system to see what clips visually on the LCD in comparison to what clips in your RAW converter. Exposing to avoid highlight clipping, at the expense of more shadow noise, is the lesser of two evils. Remember, the worst possible situation is blowing out subtle highlight data you hope to render in your image. However, the ETTR technique illustrates that underexposure causes excessive noise in shadows at all ISO settings, so this too should be avoided.

Shooting in controlled situations, where you can light the scene and use an external incident meter based on the exposure tests described and perhaps bracketing, may better lend itself to the ETTR technique. Shooting quickly in the field, where you have to rely on the camera metering and where highlight values may change at any given time, may require you to live with a bit more noise in the shadows.

Noise in shadows can be reduced in a number of ways, but if you blow out highlight data you wish to capture, there's nothing that will bring that data back. One way to reduce shadow noise is to increase the black clipping in a RAW converter, especially if there isn't any critical deep shadow detail you need to reproduce. Much of the noise will clip to black, effectively removing it from the image. This technique won't work if your goal is to render as much shadow detail as possible. Downsizing a high-resolution image also will reduce noise, since a group of adjoining dark pixels will be sampled into a single pixel value. If you can capture more than one shot of the scene, at the same exposure, you can align multiple layers in Photoshop CS3 as a Smart object and use the Median blend mode to remove noise (see http://photoshopnews.com/2007/03/27/image-stacks-in-photoshop-cs3-extended/). Lastly, there are third-party noise-reduction plug-ins available for use within Photoshop.

With respect to bracketing exposures when possible and desired, it would be useful if our cameras could bracket in one direction: from "normal" to increased exposure in a preset number of *f*-stop increments. Even slight underexposure only increases noise in the image, while the "normal" meter-based exposure always can be made darker. Actual underexposure of the data provides no benefit. It would be equally useful if the camera manufacturers would build their light meters, LCD readouts and data handling based on RAW data. Until then, it's useful to understand the role of exposing for the RAW, linear data and, if possible, capturing as much data as a RAW converter can render. You may think you overexposed an image based on an initial rendering observed in a converter when, in fact, you might have up to a full stop or more data that can be normalized using appropriate exposure compensation.

In the final analysis, ETTR isn't about overexposure, but rather proper exposure, while avoiding true highlight clipping of linear-encoded data. This often isn't the exposure our light meters recommend.