Chapter 3
Digital Exposure and Metering

No approach to editing and printing digital images will work if we can't get the image right at the time of exposure. With digital cameras, the old computing maxim applies: garbage in, garbage out (aka GIGO). This chapter will walk through the basics of metering and exposure in digital imaging, so we can minimize the G.

Metering

Most of us tend to use the reflected-light meters in our cameras. These are wonderful tools that help us to get a correct exposure, but still aren't perfect.

There are also handheld meters that can be used like in-camera meters to measure light reflected by a subject, or in incident mode, to measure the light falling on a subject. Many handheld meters are also able to measure light from studio strobes, which takes a more sensitive sensor. In-camera meters can't measure strobes.

When applying the Zone System, Adams used a handheld spot meter. Spot meters are specialized meters that can measure light reflected from a very small area, typically a one-degree arc. You can use them to make precise measurements of different areas of a scene without overlap of one area of brightness onto another, because the measurement area is so small. Most standard handheld meters do not have a spot function. If such a feature exists on a meter, it will be explicitly described by the manufacturer, and the meter will have a viewfinder that you look through to take readings.

Some in-camera meters also have a spot feature, but it's not quite as useful as a handheld spot meter. Most in-camera spot meters measure a percentage of the field of view of the lens attached to the camera. Some meter an area as small as one percent, others meter three percent, and still others meter as much as nine percent, which really doesn't make it a spot meter at all since its measuring area is so wide.

With a wide-angle lens even a one-percent spot meter can measure a wider area than desired. If your lens has a 140-degree angle of view, a one-percent spot meter will measure a 1.4-degree arc. The measurement is tied to the angle of view of the lens. Longer lenses with narrower angles of view will give you a smaller measurement area.

Most DSLR cameras have three or four metering modes, which can include a spot meter, a center-weighted average meter, an average meter, and an advanced multi-segment meter. All of these are reflected-light meters, but work in slightly different ways. We've already discussed how a spot meter works. A center-weighted average meter works by averaging the brightness across the scene as presented in the viewfinder, but giving more weight to the center of the frame. That's because a lot of people center their main subjects. Let's assume the meter has three measurement zones—one for left of center, one for right of center, and one for the center of the frame. Let's say we're shooting with an aperture of f/8. The area to the left of center is in shade, so the shutter speed meters as 1/30 second. The area to the right of center is in sunlight, so the shutter speed for that area meters as 1/500 second. The center is partly in shade, so the shutter speed for that area meters as 1/200 second. Now let's assume that the weighting is 60% for the center and 20% for each of the other metering zones. The metered shutter speed for this made-up scene would be (0.60) (1/200) + (0.20) (1/500) + (0.20) (1/30) for a weighted-average shutter speed of 1/226, right in between the two camera settings available of 1/200 and 1/250.

The average meter works by simply averaging the different areas of brightness in the scene to come up with an average brightness reading, and then setting an exposure for that average brightness. Using the same three-segment meter and scene as before we'd get a calculation of (0.33) (1/200) + (0.33) (1/500) + (0.33) (1/30) for an average shutter speed of 1/243, or the 1/250 second setting available on the camera. There isn't a large difference in this case. In others the difference may be larger.

Multi-segment meters are high-tech average meters, if you will. Modern cameras have light metering systems that take brightness readings from many different areas of a scene and process each of them discretely with an on-board computer. The active focus point(s) may also be taken into consideration to help determine how the weighting of the different readings will be done with the active focus point(s) being used as an area of higher importance. The computer in the camera will calculate a weighted average reading. Typically, the areas around the center of the frame or around active focus points are given more weight, and areas toward the edge of the frame are given progressively less weight. The scene will be analysed and compared to reference image information stored in the metering computer. Multi-segment meters are generally very accurate and do come up with a good exposure setting in most circumstances. Like all reflected-light meters, however, they can be fooled.

The 18% Gray Phenomenon

Fooled? How? Light meters are calibrated to a middle gray tone. Middle gray happens to be Zone V in the Zone System and represents an amount of 18% of the light that hits the subject being reflected back to the meter. That's why we have 18% gray cards for taking meter readings.

This means that the reflected-light meter wants to turn everything it "sees" into a middle gray tone. If you take a photo of a black cat in a coal bin, what color do you think the cat will come out as? If you take a picture of a white sheet lying on snow, what color do you think it will come out as? The answer is the same in both cases: middle gray. The meter is fooled by overly light or overly dark subjects and comes up with the wrong exposure as a result.

Figures 3.1, 3.2, and 3.3 illustrate this. The images show a hibiscus flower in my garden. The petals are white with contrasting dark pink areas. In the first image, I took the shot using my camera's spot meter, measuring off the white petal with no adjustment to the metered exposure. The camera was a Canon EOS 5D, which has a 3.5% spot meter. I used a Sigma 105 mm f/2.8 macro lens with a 23-degree angle of view, so my spot meter was measuring a 0.8-degree arc. Given this small measuring arc and how close I was to the flower—about one foot—I was able to use the spot meter and still isolate the white of the petal.

The white petal is middle gray. The histogram for the image shows us the same thing. The empty space on the right shows a lack of tone in the highlights. The peak in the middle shows us that most of the image (the petal) is middle toned, but the distribution in tones in this image is skewed far to the left, indicating an overly dark image. The in-camera meter turned the white-petaled flower into a gray-petaled flower.

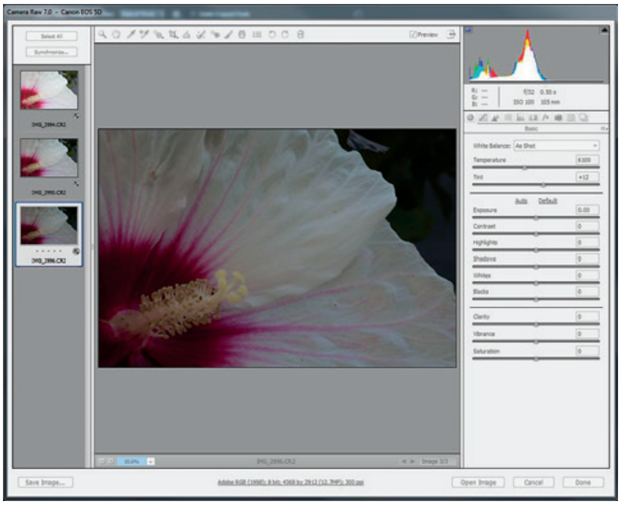

◀ *Figure 3.1: Reflected spot meter reading off the white of the petal. The image is two full stops underexposed.*

What happens if we increase the metered exposure by one stop? Having done so, figure 3.2 shows the petals are a much lighter gray, but still not white. The histogram shows that we still have underexposure, but overall less of it. The peak in the middle has moved to the right by one stop, but there is still a gap on the right tail of the graph telling us we have no whites. The distribution of tones on the histogram chart is pushed more to the right, but still not as far as it needs to be for a correct exposure.

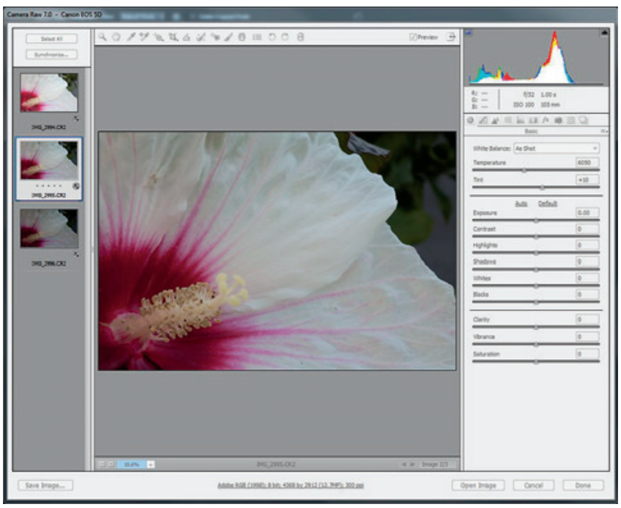

◀ *Figure 3.2: A one-stop increase in exposure lightens the petals of the flower. This shot is now one stop underexposed.*

Now let's take a look at what happens when we increase the metered exposure by two stops. As we see in figure 3.3, we now have nice white petals on the hibiscus. Just what we want! In the histogram, the peak has now moved nearly all the way to the right and there is no gap, indicating we now have white highlights. We have effectively compensated for the meter being fooled by the bright subject matter. Take note of the direction we moved the exposure in this case. The meter wanted to *underexpose* our white subject, so we had to *increase* exposure to compensate. The opposite would be true in the case of an overly dark subject. In that case, the meter would want to *overexpose* the subject, making it lighter than reality, so to compensate we would need to *decrease* exposure to darken the subject back to what it should be. With time and practice, these types of adjustments become second nature.

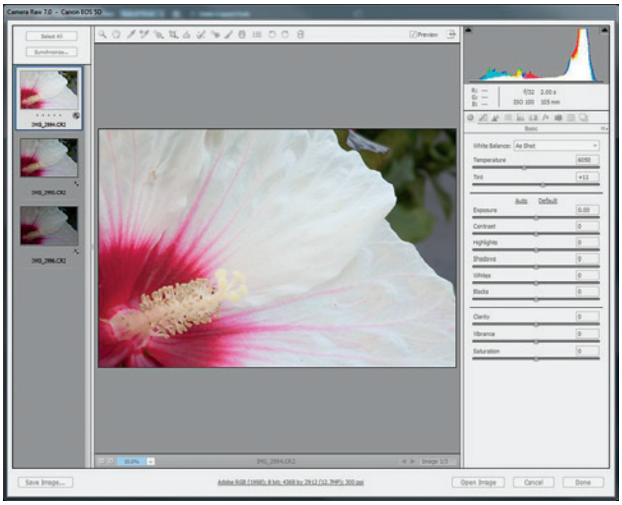

▶ *Figure 3.3: Properly exposed image showing true white in the petals*

This is how meters can get fooled. Even with very smart multi-segment meters, and with overly dark or overly light subjects, the meter will try to turn them into a middle gray. It does this with everything, actually, but the effect is most seen on very light or very dark subjects. In these cases, we have to compensate for the meter and adjust our exposure accordingly.

Incident meters—those that read light falling on the subject rather than light reflected by the subject—can't be fooled in the same way as reflected-light meters. With an incident meter, you just set your exposure at the meter reading and you're good to go. However, it's difficult to measure the light falling on a mountain from half a mile away, so using an incident meter is not always practical.

Even though reflected-light meters can be fooled, using them to measure the brightness range in a scene is still helpful. Just remember that the meter will want to turn everything into middle gray. The difference in exposure readings to turn a light subject to gray and a dark subject to gray will still be captured by the meter.

The exposure for our gray-petaled hibiscus was 1/2 second, at f/32 and ISO 100. The meter reading for the properly exposed, white-petaled hibiscus was 2 seconds, at f/32 and ISO 100. Knowing that the petal should be white and that the meter will try to turn it middle gray, we can compensate by the necessary factor to get a proper exposure the first time. We would increase the exposure by two stops from the meter reading. Recall from chapter 1 that this is essentially the "place" part of place and fall.

Test the meter in your camera under different conditions and in different situations. Figure out how each metering mode exposes the scene differently. This will help you in the field to know when you can accept the meter reading or when you have to help the camera out by adding real intelligence to the situation.

What is a Histogram?

Simply put, a histogram is a graph showing the distribution of the pixels in an image, from black on the left to white on the right.

What isn't a histogram?

As strange as it may sound, a histogram is not a substitute for a light meter.

Everything we see in the histograms in figures 3.1–3 and the assertions about whether the images are under- or overexposed is based on the context of the images.

A histogram that's skewed to the left does not absolutely mean an image is underexposed. Conversely, a histogram skewed to the right does not absolutely mean the image is overexposed. In the previous examples, the histograms showed underexposure in the context of the image. Let's look at some additional examples to see how the context of the image is important.

The histogram of the image in figure 3.4 is skewed to the right. But this is what's called a *high-key* image. A high-key image is one in which most of the tones are light and there are few dark tones. In the context of this image, the histogram is fine.

The opposite of a high-key image is a *low-key* image. In a low-key image, most of the tones are dark and there are few light tones. The histogram for this type of image is skewed to the left, as shown in figure 3.5. In the context of the image, the histogram is correct. In another context it might show underexposure. But note that the graph continues nearly all the way to the right edge. In a truly underexposed image, this would not be the case.

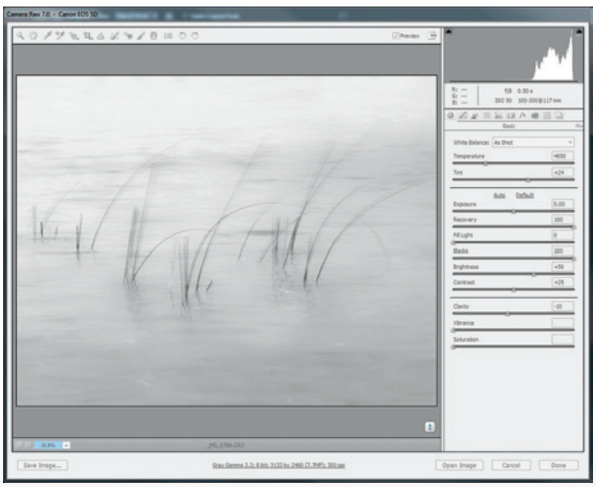

▲ *Figure 3.4: High-key image*

▲ *Figure 3.5: Low-key image*

The image in figure 3.6 looks flat and lacking in contrast, and this is confirmed by the histogram. The graph is pulled in from both ends, indicating there are no true whites or true blacks. This is a flat- or low-contrast image.

Many people feel that this shape of histogram indicates a good image. It may, but not necessarily. What it shows is that the majority of the tones are in the middle zones. This may be fine, depending on the context of the image. The histogram that approximates what in math is referred to as a "normal distribution" is neither good nor bad by default. It depends on the image. It is rare, however, that we'll want the histogram for a normal-contrast image to show gaps at either or both ends.

The last kind of image we'll look at is a high-contrast one. The histogram for a high-contrast shot fills the entire graph. Often there are peaks at the left and right ends and flatness, or a trough, in the middle. This tells us that there are true blacks and lots of other dark tones, along with true whites and plenty of other light tones. That is exactly what the graph for the image in figure 3.7 shows, and in the context of the image, it is correct. At least it's correct for me. Your vision for a similar image may be different, and that's fine.

The histogram in and of itself does not tell us whether our image is properly exposed, and it should not be used as a proxy for a light meter. Context tells the tale. What may be an incorrect exposure in one image will be correct in another. Keep this in mind when you're viewing your images in the field and when you're editing them in the digital darkroom.

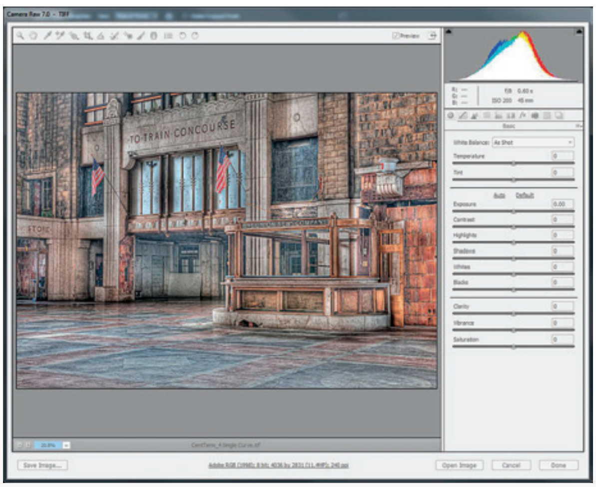

▲ *Figure 3.6: Low-contrast image*

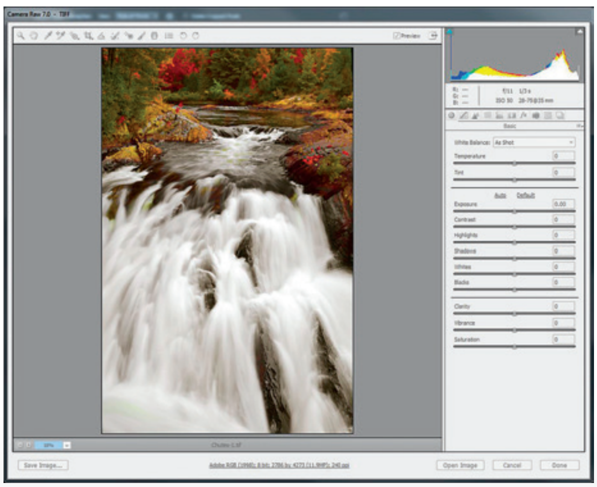

▲ *Figure 3.7: High-contrast image*

Exposure

Now that we know how light meters work, we can move along to looking at how exposure is determined and what happens when we adjust one or more of the variables that go into exposure. What I like to call the "exposure triumvirate" is made up of *shutter speed, aperture*, and *ISO*.

The exposure triumvirate is made up of shutter speed, aperture, and ISO.

Shutter speed is the amount of time the shutter stays open to let light pass through the lens and onto the sensor or film. Shutter speeds can range from fractions of a second to many seconds or minutes, and in some cases hours. For most DSLRs, the range of shutter speeds in one of the programmed exposure modes (more on these later) is from 1/8000 second to 30 seconds. When you put the camera into what's known as Bulb mode, shutter speeds of longer than 30 seconds can be obtained by holding the shutter open manually or, preferably, with a locking cable release.

The aperture is the opening in the lens that allows light through. It can be adjusted to different diameters. In most cameras, the aperture is in the lens, and the shutter is in the camera body. In some medium- and large-format cameras, the shutter and aperture are both in the lens.

Aperture settings are described by f-numbers such as f/4, f/8, f/16, and so forth. An f-number tells you the ratio of the focal length of the lens to the effective aperture, which is the diameter of the beam of light coming through the aperture. The "f" refers to the focal length of the lens. Huh?

Take a lens with a focal length of 100 mm with the aperture set to f/4. The ratio of 100:4 can also be expressed as 100/4, or 25. So the effective aperture is 25 mm. If you set the f-number to f/8, the effective aperture is 12.5 mm. In the case of a 200 mm lens with the f-number set to f/4, the effective aperture is 50 mm. Huh?

Remember that the f-number is representing a ratio. The ratio is constant, but that means the size of the opening has to vary. With longer focal length lenses, the aperture has to be larger to allow the same amount of light through.

This is important. While the size of the opening differs depending on the focal length of the lens, the amount of light that gets through onto the sensor or film is the same at any given aperture setting. So f/4 on a 14 mm lens lets the same amount of light through as f/4 on a 600 mm lens at the same shutter speed. If this weren't the case, then light meters would be useless. We need to know that the relationship between the three elements of the exposure triumvirate is the same for all of our lenses.

The ISO setting is the wild card in the triumvirate. It's not something that we would typically change unless we need to, or unless we want a certain look that comes from grain. In digital imaging, we don't have grain, of course; we have noise. Noise and grain aren't the same, and while grain can enhance some images, it is rare for noise to do so.

With film, moving from an ISO 100 film to an ISO 200 film actually meant the film was more sensitive to light. The tiny crystals of light-sensitive silver halide embedded in the film emulsion were larger on an ISO 200 film than on an ISO 100 film. As you increased the ISO rating, the grains become larger still as you moved up the scale to ISO 400, 800, and 1600 films. Those larger, more light-sensitive crystals are what made the higher-rated film "faster," or able to record an image at a faster shutter speed or smaller aperture.

With a digital camera, increasing the ISO setting does not actually make the sensor more sensitive to light. The pixel sites on a digital camera sensor have a fixed sensitivity to light. When we increase the ISO setting on a digital camera, what we're doing is boosting the gain.

In a digital camera, light hits the pixels and is recorded as a number of photons that are captured by the pixel. This analog number of photons gets translated into a voltage that is then converted to digital information, and stored with all the information from the other pixels on the memory card. The number of photons each pixel can capture is fixed. Increasing the ISO setting does not mean more photons can be captured. So what does increasing the ISO setting do?

When we increase the ISO setting, the digital signal that's captured is boosted. This comes at a cost, however. That cost is noise (the digital equivalent to grain). Noise is present in all digital images, even those taken at low ISO settings. When the signal is amplified, noise is also amplified. That's why we see noise as more evident in higher-ISO digital images than in lower-ISO images. Advances in converting the analog photons into a digital signal and in processing that digital signal in-camera have given us newer cameras with lower overall noise, which means lower noise at higher ISO settings as well.

All digital camera sensors have what is called a *base* or *natural* ISO setting. It's typically the lowest ISO setting the camera can be set at in its regular range (i.e., without going into the L/Low or H/High ISO settings). In many Canon cameras, for example, this is ISO 100. In some Nikons it's ISO 200 and in others it's 100.

Let's take a look at each component of the exposure triumvirate in more detail.

Shutter Speed

As we noted earlier, the shutter speed is the amount of time the shutter remains open to let light pass through the lens and onto the film or sensor.

The shutter speed helps us control the appearance of action or movement. Faster shutter speeds help us stop action and slower shutter speeds give blur or motion to action.

With a static subject, such as architecture, the shutter speed is less important and can be set to anything that produces a good exposure with the desired aperture and ISO settings.

When we have a subject or scene with movement, the choice of shutter speed becomes more important. Do we want to freeze the movement or allow it to blur slightly to give a sense of motion? It depends. Sometimes freezing movement is the way to go and sometimes not. In many sports images we see sharp, frozen images. The action depicted by the subject lets us know there's movement going on. But some of the more compelling sports shots have motion blur in them. I recall a shot like that taken from above the bobsled track in Vancouver during the 2010 Winter Olympics. The photographer used a longer shutter speed, and the sled was a blur of color going through a curve in the track. It was not a traditional sports image, but it was incredibly compelling nonetheless. Sometimes it can be preferable to allow just a bit of motion blur into the shot. A race car on the track, for example, can be given a sense of dynamism with a little blur in the wheels.