

# The Fotobug's Guide to the Histogram

by Jim Caldwell


## The Fotobug's Guide to the Histogram

*M*any digital cameras, except perhaps the low-end point-and-shoots, offer a tool that many new photographers, and perhaps a few seasoned ones, don't really understand. I'm referring to **The Histogram**.

OK, so what is a histogram? Good question, glad you asked! A search on the web defines a histogram as “a diagram consisting of rectangles whose area is proportional to the frequency of a variable and whose width is equal to the class interval”.

Say what?


Well, let's simplify this a bit and narrow our definition to the use of a histogram for photographers. Essentially, a histogram is simply a graph of the luminance values in a digital image. Luminance values are the numbers that represent the darkness and lightness of the pixels on our images. A basic histogram looks like this:



Now, the actual histogram that is displayed on your particular camera may vary somewhat from this image, but just follow along since regardless of how it actually looks, it will still represent the same thing!

Each pixel luminance value is graphed on this chart of the image that is displayed on the LCD of your camera. The height of the graph represents the number of pixels that contain that particular value. The left side of the chart are the darkest values in your image and the right side are the lightest values, like this:

Dark Side  
(shadows)



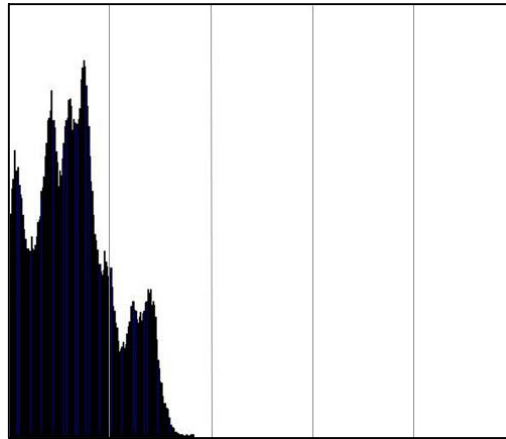
Light Side  
(highlights)

So, the darkest pixels in your image are represented by the curve on the left of the chart and your brightest pixels are on the right and the midtones are right in the middle. Each of the vertical lines represents an approximate one stop range, so in this case the entire chart represents five stops from the left side to the right side. In the chart above, the graph represents a fairly even distribution of luminance values, with a little peak in the left side (the first stop) and a lesser one in the right. I would suspect the image that this is based upon would contain a nice even range of shadows to highlights with most of the tonal values falling right in the middle.

Unfortunately, our sensors are not yet capable of capturing the entire dynamic range of light values found in nature without resorting to techniques such as HDR (High Dynamic Range) which involves taking a series of images at various exposures. So, as a photographer, our job is to try to squeeze as much of that image detail from the scenery in front of us onto the sensor without blocking up the shadow detail or blowing the highlights.


Let me explain ‘blocking’ and ‘clipping’ a bit further.

If we underexpose an image, that means that the graph as seen above will be all the way over on the left side of the graph and, in fact will be flattened against the left side, like this:



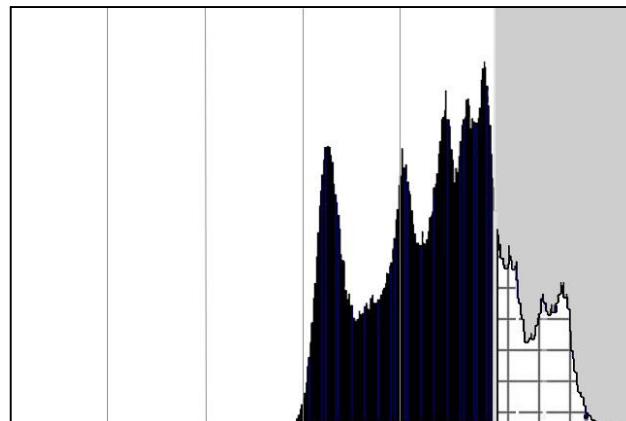
This image is at least three stops underexposed (note the graph of values stops in the second vertical box and the right three have nothing in them!). Also, the left edge of the graph is flattened up against the left side. This means that the sensor has only captured about a two stop range of values and shadow detail has been lost. If you would try to correct this image in your favorite image program on your computer, the final image would contain very muddy shadows, would be quite grainy from the boosted sensor noise values, and would contain little image detail. All of that detail has been lost forever.

Let's take a look at the other side of the coin. This histogram represents an image that is severely over-exposed:



Note that the graph is now all the way over on the right side and is flattened up on the right edge of the chart. Everything that falls outside of the chart has been 'clipped' – those values are beyond the capacity of the file, or sensor to record them and the lighter values will be 'blown out'. Also note that there are no values appearing in the first two vertical boxes and just barely on the right edge of the third box from the left.

Once again, trying to "fix" this image on your computer will result in a final high contrast image with extreme highlights that contain no detail. That detail has fallen over the right side and is gone forever. In the following image, I show how only a portion of the luminance from the image was captured by the graph above:



Everything over on the right in the grey area was outside the capture range and once gone, it is gone for good. In this instance, reducing the exposure by about three stops would bring it back into the histogram, safe and sound!

Let's review what we have learned so far. The histogram is a chart of the luminance (light to dark) values in our image. In order to capture the maximum amount of detail, we want to make sure that the values lie within the chart and don't fall outside on either the left or the right. If the chart shows the graph bunched up on the left, the image is underexposed and we need to increase our exposure (open our aperture, decrease shutter speed, or raise the ISO). If the graph is all the way over on the right, our image is overexposed and we need to reduce exposure (close the aperture, increase shutter speed, or lower ISO).


Don't get cocky, now – there is still more to learn!

Consider the following image I took of a timber wolf at sunset:



This image has a lot of dark shadows, but the wolf looks properly exposed. Let's take a look of the histogram for this image:






As we should have expected, there is a spike on the left side of the histogram, which is all those dark values in the wolf image. However, note that the values still go all the way across the chart to the right side, and then stop. There are no values that appear to go beyond the right side, although some values may have fallen off on the left. However, this histogram shows that his particular image is properly exposed. Ignore the height of the chart values, we are only interested in the range from the left to the right. If I increased the exposure on the wolf, we would have picked up more detail in the shadows, but at the expense of the highlights on the fur and the eyes. Since the wolf is our subject, we don't want him to be improperly exposed! So, screw the shadows!

Please note that in a case such as this one, the graph may be difficult to read. The values of the mid-tones and highlights may be so few that you simply can't see that part of the chart across the bottom edge on the back of your camera. That's why it is important to evaluate the histogram as it relates to the image you are capturing! If you are in doubt, change your exposure to bracket the image and take another (if possible!). If the image appears to be a bit "foggy" or low-contrast on your LCD, but the histogram appears to have values all the way over to the right side (and NOT beyond!), then the image is likely OK and once you have applied the necessary minor curves and/or exposure adjustments to that image in your computer, it should look fine!

How about one more example?



No, that isn't snow, it's mineral deposits in Turkey. However, this example could just as easily apply to a snow scene, or a beach scene with pure white sand, etc. Here is the histogram for this image:



Just as we should suspect by now, there is a peak over on the right side where the lighter values are found (again, the height just indicates how many pixels are in those areas) and the histogram contains values all the way over to the left. Also, the highlights go right up to the right side, but not beyond. This is a properly exposed image.

Before I wrap up this discussion, let me try to explain why it is important to capture values across the range of your histogram, especially to the right side edge.

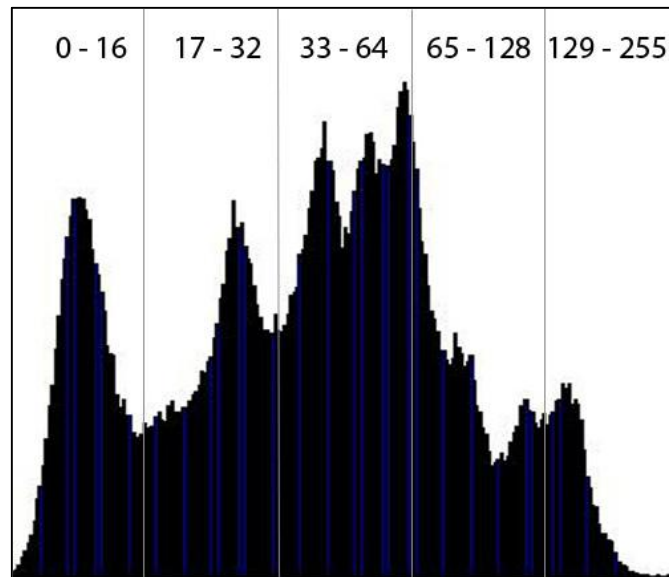
Digital devices, such as our cameras, record values as binary numbers, simple ones and zeros. A bit is on or it is off, there is no in-between. So, we allow a bit value that is on to represent 1 and a bit that is off represents 0. A series of these bits can be used to represent larger numbers. A JPG image uses eight bits to represent a number so the range can be 00000000 to 11111111. The first example is (of course) zero and the second value would be the binary, eight-bit representation of 255. Now, in a real JPG file, the final pixel value is actually based on three of these eight bit numbers, one for red, one for green, and one for blue; but for the purposes of this discussion, let's just assume a single value, zero to 255. In this case, zero is pure black and 255 is pure white, with a range of grey values in between. The largest number that can be represented is 255. If our sensor tries to record a value beyond 255, 255 is all that will be retained. The file simply cannot hold a number bigger than 255! Any value over 255 would be those values that are clipped and lost in our final captured image – forever! You really can't put 10 pounds of fertilizer in a 5 pound bag!

So, 255 is the maximum! This would be the brightest white that we can record for our image. Now, hopefully you will recall that the vertical lines in the histogram represent

one stop of exposure. If 255 is the value at the very right edge, 255 is our brightest value, then the first line to the left of this would be one half of that value since a one stop down value would be half the amount of light (each stop up represents twice as much light and each stop down one half!). So, half of 255 is 127.5, but lets just round it off and call it 128. Ok, follow me closely on this! 255 is the right edge and the first line (stop) to the left of that is 128. So any recorded value between 129 and 255 will lie on the chart between the right edge and that first line.

Continuing on to the next line to the left (2 stops from the right edge), the value would be halved again so one half of 128 is 64. Therefore, any recorded value from 65 to 128 will fall between the two lines that are the first and second lines from the right.

We can continue on all the way to the left side of the chart which is our zero value. However, stop and notice something interesting. Each range of those one stop values contains an unequal range of values! The first stop down from the right contains a range of 127 values! That is the same as the other four stops combined!



This is due to the linear nature of digital capture. Each one of those values represents detail in our final image, so if we have underexposed our image by only 1 stop, we have lost half the detail that could possibly be captured! That is why a well-known landscape photographer coined the phrase “expose to the right”. That means that you want that histogram to go as far to the right side as possible, without going beyond.

Now, to further complicate this discussion, I have been basing the examples on an eight bit captured image, such as a JPEG image. If you are shooting RAW, many cameras are capable of capturing up to 14 bits (perhaps even more). There is a very good chance that the histogram representation on your camera is still based upon the eight-bit image (RAW embeds an eight bit thumbnail in the data that it stores) and since 14 bits can contain a much higher number (16383 to be exact!), if you are capturing in RAW, and if



the histogram shows that you are slightly over the right edge, there likely isn't a problem since the RAW capture has a broader range. Just one more reason to be shooting RAW!

Some cameras are also capable of showing a histogram that represents the three color values (red, green and blue) as well as the average luminance values. It is possible that the average value falls within the graph, but one of the color values actually goes beyond the edge. This normally isn't a problem and not one to get hung up on. I rarely use the color histogram on my cameras.

Finally, remember to evaluate your histogram based upon the captured image! If you are shooting a dark bird against a white sky (boy do I hate that!) – remember that it is the bird and NOT the sky that is your subject (at least I would expect that to be the case!). So, if the sky is clipped at the right edge of your histogram, that is OK as long as the values that represent the bird aren't clipped. Use common sense!!

Since a histogram is a summary of the tonal values in an image, consider that a scene with a good range of tonal values should have a histogram with a nice curve from the left to the right, with most of the values falling somewhere in the middle. A high contrast situation would have a histogram with spikes on the left and right, with fewer values in the middle. A very low contrast scene should contain a narrow band of values somewhere in the middle of the histogram with a few values approaching the left and right sides. Start using the histogram and soon it will begin to make sense!



*We* hope you enjoyed this free guide and hopefully it will help you make better exposed images that capture the maximum amount of detail that your camera(s) are capable of capturing. If you have any questions, don't hesitate to contact us and please continue to tune in to our bi-weekly podcasts at [www.thefotobug.com](http://www.thefotobug.com) - or better yet, via one of the major podcast feeds such as iTunes, Zune, or Miro! You can also return the favor to us by leaving us some feedback on iTunes, (or Miro, or even at our website, [www.thefotobug.com](http://www.thefotobug.com)). We are also on Facebook and would love for you to 'Like' us!

There is more to come!

© 2011 Jim Caldwell

Authored by Jim Caldwell –  
[www.thefotobug.com](http://www.thefotobug.com)  
[jim.fotobug@verizon.net](mailto:jim.fotobug@verizon.net)