The Noise about Noise

I have found that few topics in astrophotography cause as much confusion as noise and proper exposure. In this column I will attempt to present some of the theory that goes into determining the “correct” exposure for a given scene and then show some simple guidelines that can make it easy – at least for the DSLR users. As someone that works with a variety of signal processing systems in my day job, I've been well acquainted with noise and its properties. Noise, by its random nature can be confusing, but with a little knowledge we can quiet most of the *noise* about noise and take steps to control its effects in our images. The whole idea is to figure out the proper exposure to reduce noise as much as possible and produce good quality data ready for processing.

What is Noise

First let’s get a working definition of noise as it relates to imaging. Officially, noise is any artifact in an image that is not present in the actual scene. For processing purposes, this is a little broad, as it would encompass any optical defects as well. Generally, noise is a random image artifact that is a function of a component in the data acquisition system or a function of the scene itself. In this case, the former means the camera excluding the optical system (scope or lens) and the latter means photon noise. As we will see, this random aspect of noise is very important in combating its effects.

Types of Noise Encountered in Astrophotography

There are generally two noise sources we are concerned about in astrophotography. Broadly the two categories are photon or image noise and camera noise.

Dark current noise, quantization noise and read noise are the usual culprits for camera noise Dark current noise is the one with which we are most familiar; it is the signal that builds up in the sensor even without any light falling on the chip. This noise is proportional to both the exposure time and the temperature. Dark current noise can be modeled as a combination of a fixed, deterministic value that is dependent on temperature and exposure time, and a random variation with a zero mean about this fixed value. In fact, it is because part of the dark signal is constant that we can remove it with a dark frame. If we look at just one pixel in a dark frame and plot its values over many images, we get a curve that looks like the one shown in Figure 1.

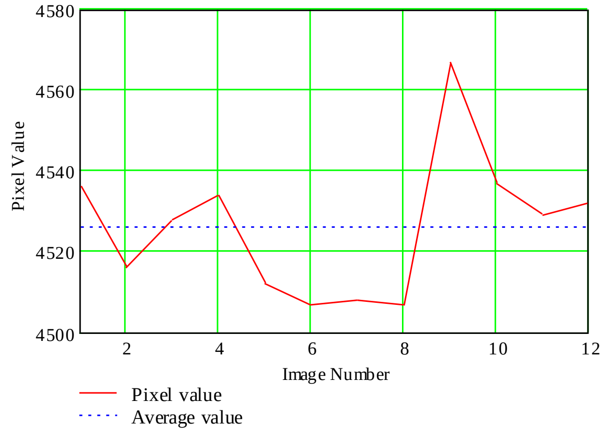

Figure 1 - Single pixel value over several images

If we examine the plot in Figure 1, we see that the pixel has an average value (blue dashed line) of 4526 and a random fluctuation around that average. If we produce a similar plot for every pixel in the dark frame, we would find that each one has its own average value. A well-averaged dark frame is an image made up of these constant values. Subtracting a dark frame removes the average dark signal from the image producing a much less noisy picture. What remains is the actual image data plus the random variation in the dark signal collected during the exposure time. It is this random fluctuation that makes up the remaining dark current noise and it can be thought of as a zero-mean random signal. This remaining noise is proportional to the square root of the integrated dark current; it scales with temperature such that it doubles roughly every seven degrees.

Readout noise is caused by noise in the analog amplifier chain between the sensor and the analog to digital converter or ADC. This noise is fixed in level and unlike dark current noise is not proportional to exposure time.

Quantization noise results from the fact that the ADC outputs only discrete integer values. If the actual data falls in-between possible ADC values, an error or noise results.

For purposes of this discussion we will ignore the effects of quantization noise, as it is small compared to the others in modern cameras with 12 to 16 bit ADC's. Instead we will consider only read and dark-current noise.

The remaining source of noise comes from the image itself. Quantum mechanics tells us that light itself is noisy; photon noise is inherent in light and has a Poisson distribution with an average value equal to the square root of the number of photons collected at each pixel.

The SNR Equation

There is a classic equation describing the signal to noise ratio of the data collected at each pixel in a digital camera: $\mathrm{SNR} = \sqrt{n}\cdot\dfrac{s_{obj}\cdot t}{\sqrt{s_{obj}\cdot t + s_{sky}\cdot t + s_{dark} + n_{read}^{2}}}$

Where:

n = number of sub-exposures and assumes an "average combine" in the final image.

S_{obj} = Object flux in electrons per unit time

S_{sky} = Sky background flux in electrons per unit time

t = Exposure time in units matching the units used for S_{sky} and S_{obj}

s_{dark} = Number of dark current electrons

n_{read} = Readout noise in electrons

From this we can readily see that the SNR improves with the square root of the number of sub-exposures. With a little more inspection we can also see that it improves with the square root of the exposure time. Now comes the interesting part. From the definition of SNR we have $SNR = \frac{signa}{noise}$ and the pixel SNR equation then tells us that the signal is $s_{obj} \cdot t$ and the noise is $\sqrt{s_{obj} \cdot t + s_{sky} \cdot t + s_{dark} + n_{read}^2}$. The expression for the noise breaks down into two terms: image noise made up of $s_{obj} \cdot t$ and $s_{sky} \cdot t$ plus camera noise made up of s_{dark} and n_{read}^2. Now if the image noise is much larger than the camera noise, we can ignore its effects and the pixel SNR simply becomes $SNR = \sqrt{n} \cdot \frac{s_{obj} \cdot t}{\sqrt{s_{obj} \cdot t + s_{sky} \cdot t}}$. This tells us that if we can make the image noise much larger than the camera noise then using n exposures of t seconds is identical to a single exposure of n times t seconds, assuming the short exposures are averaged. This conclusion is of great interest in astrophotography, because it is much easier to take multiple short exposures than a single long one. If something goes wrong, you lose a single short exposure rather than the whole thing! This all boils down to one question – how do we insure that the image noise is much greater than the camera noise?

The Noise Myth

The first thing we need to understand is that it is not necessary to keep noise to low values in our data. The absolute level of the noise, provided it does not cause saturation of the electronics or the image file format, is meaningless. It is only the ratio of the image signal to the noise that matters; everything else can be scaled and manipulated in your image processor. To demonstrate this point the following simulated star images were generated using mathematical modelling software.

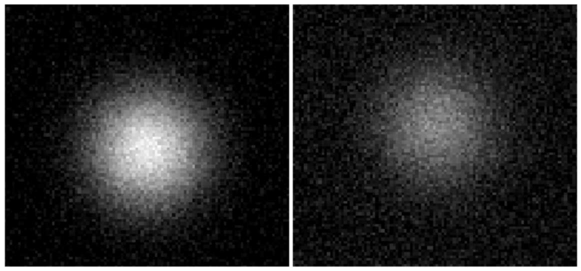

Figure 2 - Simulated star at different SNR's

The image on the left actually has a higher noise level than the image on the right, but because it has a higher SNR, it looks much better. Both images have been stretched in the same fashion to make the noise obvious.

Determining the "correct" exposure

The basic problem here is how to determine the correct exposure for a given image. First off I'd like to point out that there is no one correct exposure. Like all photography, this all depends on what part of the scene you are trying to capture. Many objects have a wide range in brightness, and you may want to choose a short exposure to better capture detail in the bright areas. The definition I'm using here is to give the best SNR over the whole of the image, even if it allows the brightest parts of the scene to saturate.

Our goal is to determine what sub-exposure will allow the image noise to dominate the camera noise and let us safely ignore the effects of camera generated noise. This is the very definition of a sky-limited exposure.

The first step in this exercise is to see how camera noise and image noise combine. There is a branch of mathematics (if you think of statistics as mathematics) that shows us that the average value of the summation of two or more random sequences is equal to the square root of the sums of the squares of the average value of the individual sequence. Using this relationship, we can examin how the total noise varies as the ratio between the image noise and camera noise changes. If we plot the ratio of total noise to image or sky noise against the ratio of image noise to camera noise we get the following plot.

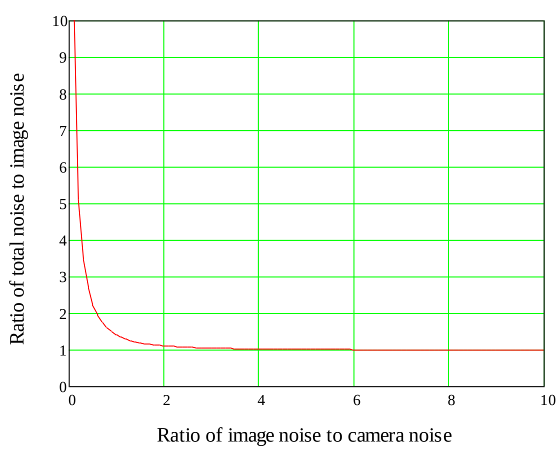

Figure 3 - Total noise ratio verses image noise ratio

As you can see from the plot, shortly after the image noise increases to twice the value of the camera noise, then for all intents and purposes the total noise is simply the image noise. Now if we zoom in on the area of the plot around the inflection point we can see things in better detail.

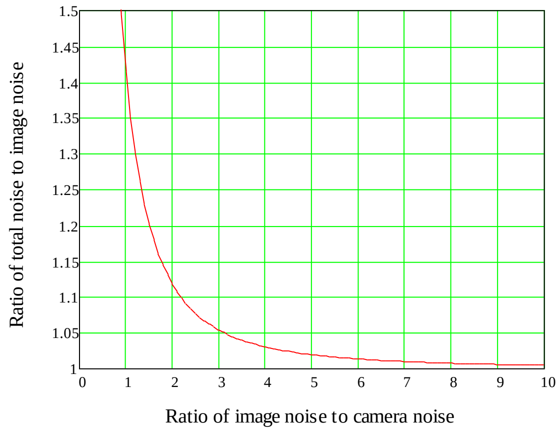

Figure 4 - Total noise ratio verses image noise ratio

An image is generally accepted as sky-limited if the total noise increases by no more than five percent due to the addition of camera noise. Using the plot in Figure 4, you can see that this occurs when the sky or image noise is approximately three times the camera noise. If we are averaging sub-exposures, then exposing each sub beyond this limit is of little value as seen from the pixel SNR equation.

So now we have a working definition of a sky-limited exposure; simply expose each sub until the sky noise is three times the camera noise. Now the problem becomes one of determining just how long an interval this is. To do this, we have to revisit the pixel SNR equation and make a couple of assumptions. Firstly we must assume that all the sub-exposures have been properly calibrated using a well-averaged dark frame. Secondly we assume that the sky signal is greater than the object flux, the case in most astrophotography. This means that the contribution of dark current to noise is greatly reduced and the SNR equation simplifies to $SNR = \sqrt{n} \cdot \frac{s_{obj} \cdot t}{\sqrt{s_{sky} \cdot t + n_{read}^2}}$. Here the background noise depends on the sky level assuming that it overwhelms the camera read noise.

Now we need to be able to measure the sky level and to do this, you need to know the system gain of your camera in terms of electrons per ADU (analog to digital converter units). You can find this in your camera manual, on the Web or you can measure it directly. Once you know this value, you can use a test exposure to measure the sky background. Take a short exposure, in which the sky is well below saturation, but where the histogram is completely separated from the left side of the plot. The number of electrons captured is calculated as follows $s_{sky} = \frac{gain \cdot ADU}{t_{exposure}}$ where s_{sky} is the sky flux and $t_{exposure}$ is the exposure time. Now remember the goal here is to make the sky noise three times the camera noise, so knowing the sky flux and the fact that noise adds as the square root of the sums of the squares, we just need to find the camera read noise and we can calculate the required exposure time. The read noise, like the gain, can be obtained from your camera manual, the Web or it can be measured. Finally the exposure time can be calculated from $3 \cdot n_{read} = \sqrt{s_{sky} \cdot t}$. After which, solving for t we obtain $\frac{9 \cdot n_{read}^{2}}{s_{sky}}$. This is the exposure time required to make the image noise of the sky background equal to three times the camera read noise. This method gives us an accurate exposure time, but it is a bit of a pain to do each time you go imaging. It turns out that there is a short cut that uses the above method for calibration. Use the above method to determine the sky-limited exposure then take an exposure using the calculated time. Examine the histogram of this sky-limited exposure and note where the peak of the histogram is located.

The next time you want to know the sky-limited exposure time for any given conditions, take a test exposure and note where the peak of the histogram is located. Then simply figure out how much more or less exposure time is needed to move it to the position found above to obtain a sky- limited exposure. I've calibrated three Canon DSLR's using this technique and in each one, the sky-limited position of the histogram was one quarter of the way from the left hand side of the histogram plot.

Let's take a look at an example using my Canon 60Da. Suppose a test exposure of two minutes produces a histogram with a peak at the one-eighth point. Since CCD and CMOS sensors have a linear response with integration time, the exposure should be increased to four minutes to produce a sky-limited sub-exposure. That's all it takes; a calibration session to know where to place the histogram peak and a simple test exposure when you go imaging.

How many sub-exposures

Now let's go back to the pixel SNR equation. We notice that the final-image SNR scales with the square root of the number of sub-exposures. This means that each time the number of subs is doubled the SNR increases by a factor of 1.414 as shown in Figures 5 through 9.

Figure 5 - Single 5 minute exposure

Figure 6 - Average of two 5 minute exposures

Figure 7 - Average of four 5 minute exposures

Figure 8 - Average of eight 5 minute exposures

Figure 9 - Average of sixteen 5 minute exposures

As you can see from the above images the SNR improves each time the number of subs is doubled, but visually the improvement from eight to sixteen images is less apparent than between one and two. Even though the SNR has improved by the square root two at each doubling, the noise becomes smaller compared to the signal as the number of images is increased, and so the eye begins to lose the ability to distinguish the difference. We can determine how the SNR improves as the number of subs increases mathematically, but this doesn't really tell us much, as it does not take into account the way the human eye perceives changing SNR.

What you consider as a sufficient number of subs depends heavily on the imaging conditions and your setup. If, like me, you have to lug your equipment to a dark-sky site, then an hour or two of imaging is usually all you can achieve in one session. If you have a more permanent installation, then spending many hours on a target is not out of the question. There is a law of diminishing returns at work here; if you have collected three hours of data then six hours will offer only marginal improvement. If after three hours you are almost happy with the result then perhaps a little more noise reduction is better than another three hours of exposure.

Sometimes, especially if your imaging time is limited and you want to get several targets, it is nice to have a rough idea of how many subs are required to get a decent image. You can calculate everything you need with the help of a little integration, but I prefer to simply get a rough calibration for my optical and camera systems and use those to calculate the number of subs required.

The goal is to measure the sky brightness, calculate the target brightness, then use the SNR equation to determine the required number of sky-limited sub-exposures to achieve the desired SNR.

First, calibrate your system. This can be done anytime and does not need to be repeated each imaging session. The calibration process will relate surface brightness and integration time to ADU values in your camera. We start by taking a sky-limited test image and measure the average level of the background with your image-processing software. Determine an average value from a few places on the scene to get a more accurate result. Divide the ADU value for the background of the image by the integration time in seconds. This gives us a value of ADU per second for the energy being received by your camera through your optics. Next we have to measure the sky background brightness. You can use a sky quality meter or simply use your test image and the technique developed by Samir Kharusi at http://www.pbase.com/samirkharusi/image/37608572. Convert the result from magnitudes, which is a log scale, to a linear value by using $\text{Linear brightness} = 10^{\frac{\text{magnitude}}{-2.5}}$. Finally divide the ADU per second value obtained in the previous step by the linear sky brightness just measured to obtain a calibration value

When planning your imaging session, find the integrated magnitude of your target and its size in square arcseconds. Convert the brightness to its linear value and divide by the size to get the surface brightness of the target. Then multiply the result by the calibration value you have obtained for your system. This now tells you the number of ADUs per second you can expect from the target through your optical system. When you get ready to image, use an SQM or Samir's technique to measure the sky background. Convert the sky background to linear and multiply by the calibration constant then plug the calibrated object brightness and sky brightness into the SNR equation, $SNR = \sqrt{n} \cdot \frac{s_{obj} \cdot t}{\sqrt{s_{obj} \cdot t + s_{sky} \cdot t}}$ to calculate the sub-exposure SNR with n set to one. The last step is to figure out the number of subs required.

Generally a SNR of 36 to 40 is required for a smooth image that can take an aggressive stretch without breaking down into a blurry noisy mess. The Horsehead shot shown in Figure 9 had a SNR of 36 when all 16 subs were averaged, and before any stretching. So I'll suggest that 36 is an acceptable SNR value. Using this we can estimate the number of sky limited subs required to be $(36/\text{sub SNR})^2$.

Now all this may seem like a lot of work, but keep in mind that it is very easy to put the math in a spread sheet that can be run in something like *Documents to Go* on a smart phone. All that is required is to fill in the object magnitude, its surface area and take a quick measurement of the brightness of the night sky. Plug those values into the spreadsheet and presto you have an estimate of the number of subs and how long each one has to be for a low-noise image. I've tested this technique on several of my older images and it agrees with the measured SNR of the stacked images to within a few percent.