Noise in photographic images

Noise in photographic images

Introduction

Noise is a random variation of image density, visible as grain in film and pixel level variations in digital images. It is a key image quality factor; nearly as important as sharpness. Since it arises from basic physics— the photon nature of light and the thermal energy of heat— it will always be there. The good news is noise can be extremely low— often impreceptably low— in digital cameras, particularly DSLRs with large pixels (5 microns square or larger). But noise can get ugly in compact digital cameras with tiny pixels, especially at high ISO speeds.

In most cases noise is perceived as a degradation in quality. But some Black & White photographers like its graphic effect: Many favor 35mm 35mm Tri-X film. The pointillist painters, most notably George Seurat, created "noise" (specks of color) by hand; a task that can be accomplished in seconds today with Photoshop plugins.

But by and large, the majority of photographers, especially color and large-format photographers, dislike noise with good reason.

Noise is measured by several Imatest modules. Stepchart produces the most detailed results, but noise is also measured in Colorcheck, SFR, and Light Falloff.

Appearance

(A) Const. sensor noise	(B) Const. pixel noise	(C) No noise	(D) Canon EOS-10D ISO 1600

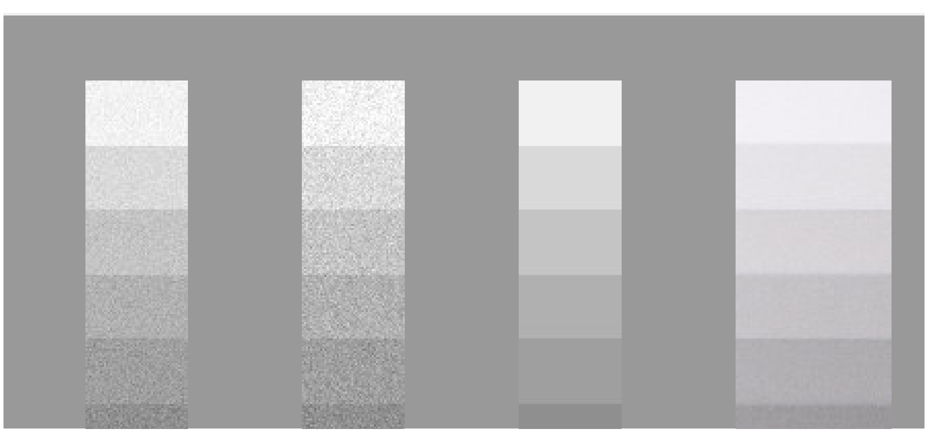

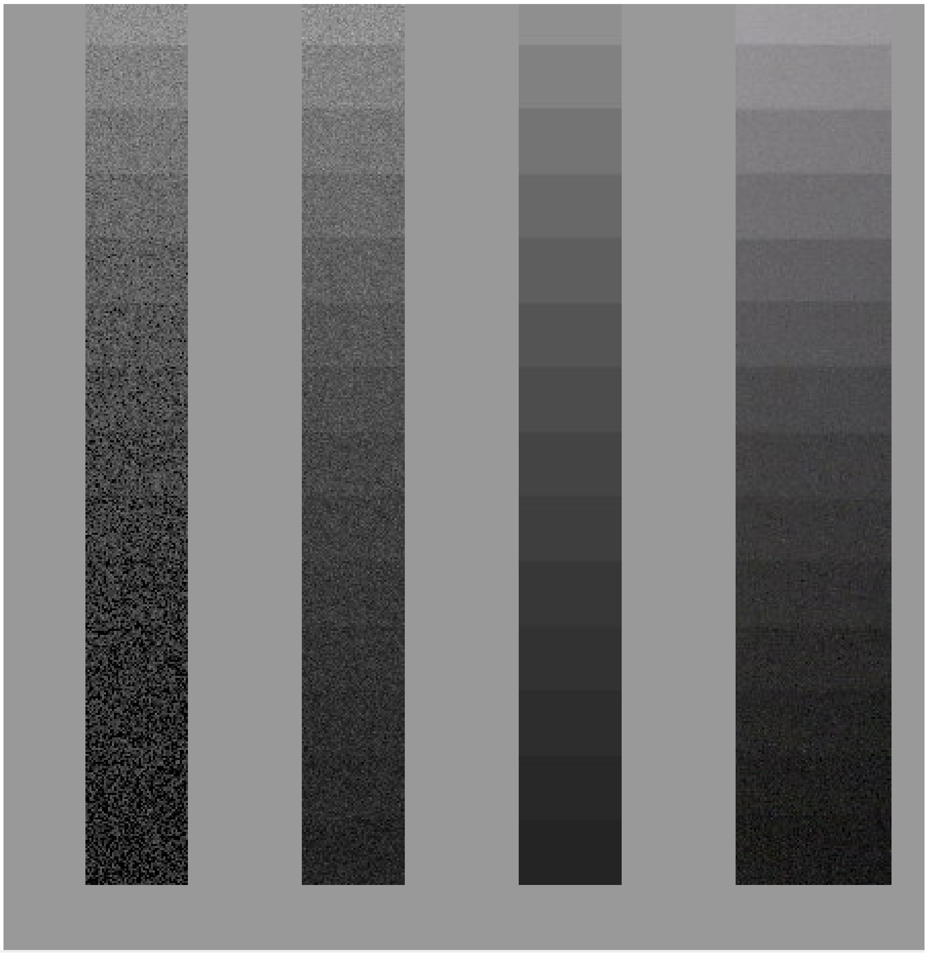

The appearance of noise is illustrated in the stepchart images on the right. Noise is usually measured as an RMS (root mean square) voltage. The mathematics of noise is presented in a green box at the [bottom of this page](#).

The stepcharts in columns (A)-(C) are simulated. They are assumed to have a minimum [density](#) of 0.05 and density steps of 0.1, identical to the Kodak Q-13 and Q-14. They have been encoded with [gamma](#) = 1/2.2 for optimum viewing at gamma = 2.2 (the Windows/Internet standard). Strong noise— more than you'd find in most digital cameras— has been added to columns (A) and (B). Column (C) is noiseless.

The fourth column (D) contains an actual Q-13 stepchart image taken with the Canon EOS-10D at ISO 1600: a very ISO high speed. Noise is visible, but admirably low for such a high ISO speed (thanks, no doubt, to software noise reduction).

The noise in (A) is constant inside the sensor, i.e., before gamma encoding. When it is encoded with gamma = 1/2.2, contrast, and hence noise, is boosted in dark areas and reduced in light areas. The Kodak publication, [CCD Image Sensor Noise Sources](#), indicates that this is not a realistic case. Sensor noise tends to increase with brightness.

The noise in (B) is uniform in the image file, i.e., its value measured in pixels is constant. This noise must therefore increase with brightness inside the sensor (prior to gamma encoding), and hence is closer to real sensor behaviour than (A). Noise appears relatively constant except for the darkest zones, where it's not clearly visible. Noise is often lower in the lightest zones, where a tonal response "S" curve superimposed on the gamma curve (or saturation) reduces contrast, hence noise.

For this reason the middle zones— where noise is most visible— are used to calculate the average noise: a single number used to characterize overall noise performance. We omit zones where the density of the original chart (hence display density in an unmanipulated image) is greater than 1.5 or less than 0.1.

Noise measurements

Noise measurements should ideally

- correlate with perceived appearance,
- (in many cases) be referenced to the original scene so the measurement is not affected by the tonal response (contrast) of the camera or RAW converter,
- be detailed enough to allow accurate assessment of sensor and camera performance, and
- be simple enough to interpret without difficulty.

Since noise is only meaningful in relationship to a signal, a ***Signal-to Noise Ratio*** (**SNR** or ***S/N***) is often calculated. ***SNR can be defined in many ways,*** depending on how Signal *S* is defined. For example, S can be an individual patch pixel level or the pixel difference corresponding to a specified scene density range (1.45 or 1.5 are often used for this purpose). ***It is important to know precisely how SNR is defined whenever it is discussed.*** SNR can be expresed as a simple ratio (*S/N*) or in decibels (dB), where SNR (dB) = 20 $\log_{10}(S/N)$. Doubling S/N corresponds to increasing SNR (dB) by 6.02 dB.

Since these requirements can be somewhat contradictory, Imatest modules have several noise and SNR measurements, some simple and some detailed. Available displays in Stepchart and Colorcheck include:

- ***Noise a function of pixel level or exposure, expressed in***
    - **Pixels, normalized** to the pixel difference corresponding to a scene density range of 1.5 for Stepchart (comparable to the density range of 1.45 for the GretagMacbeth ColorChecker) or the White – Black zones in the ColorChecker (row 3, patches 1 – 6; density range = 1.45). Without this normalization, noise is a function of RAW converter response curve; it increases with converter contrast.
    - **Pixels, normalized** to the maximum pixel value: 255 for 8/24-bit files.
    - **Pixels** (maximum of 255 for 8/24-bit files)

- **f-stops** (or EV or zones; a factor of two in exposure), i.e., *referenced to the original scene*. Noise measured in f-stops corresponds closely to human vision. See f-stop noise, below.

- ***Noise as a simple (average) number***
  - A good single number for describing noise is the average noise of the Y (Luminance) channel, measured in pixels (normalized scene density difference of 1.5), shown in the lower plot in the figures, below. The lightest and darkest zones, representing chart densities greater than 1.5 or less than 0.1 are excluded from the average.

- ***S/N or SNR (dB) as a function of pixel level or exposure,*** where *S* is the pixel level of the individual test chart patch.

- ***S/N or SNR (dB) as a simple (average) number.***

- ***a noise spectrum*** plot, described below.

In addition, Light Falloff displays a spatial map of noise.

Characteristic Stepchart results are shown below.

Column (B): Uniform pixel noise

This image has simulated uniform pixel noise (i.e., constant noise in the image file, measured in pixels).

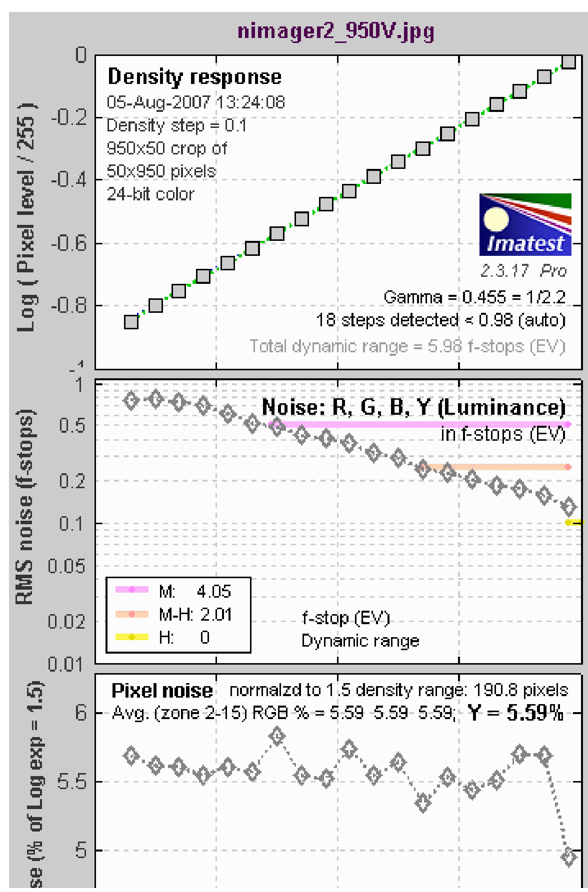

The upper plot is the tonal response (or characteristic curve) of the camera. It shows the expected ideal response for encoding with gamma = 1/2.2 = 0.4545: a straight line with slope = 0.4545 for the log-log plot.

The middle plot shows noise measured in f-stops (or EV or zones). Noise increases more steeply for dark regions (large negative values of Log Exposure) in actual sensors.

The lower plot shows noise measured in pixel levels, normalized to the difference in pixel levels for a density range of 1.5. It is relatively constant, showing only statistical variation, except for the brightest level (on the right), where the noise is reduced because some samples are clipped at pixel level 255.

The lower plot also contains the single number used to characterize overall noise performance: the average Luminance channel noise (**Y =**

5.59% , on the right near the top). This number is very high; it corresponds to poor image quality.

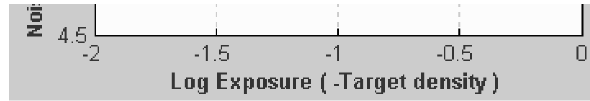

This is SNR in dB for simulated uniform pixel noise. Because of the gamma = 2.2 encoding, SNR inproves by 6.02/2.2 = 2.74 dB for each doubling of exposure (0.301 density units); roughly 9.1 dB per decade (1 density unit).

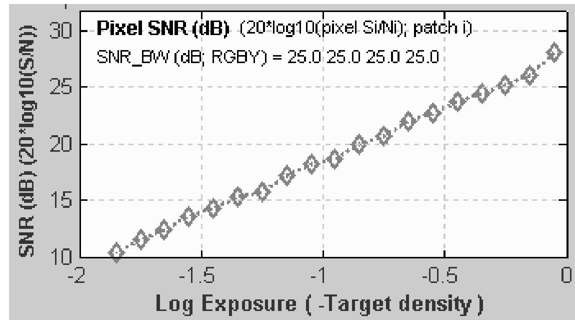

Column (D): Canon EOS-10D at ISO 1600

Unlike the above figure (for the image in column (B)), real data for the Canon EOS-10D at ISO 1600 is used.

The upper plot shows a slight tonal response “S” curve superimposed on the gamma curve (which is a straight line in this log-log plot). Some converter settings (such as “low contrast”) result in a much more pronounced “S” curve.

The middle plot shows f-stop noise, which increases dramatically in the dark regions. The lower plot shows the normalized pixel noise. It increases in the dark regions. This increase is due in part to the high ISO speed. Digital cameras achieve high ISO speed by amplifying the sensor output, which boosts noise, particularly in dark regions. This curve looks different for the minimum ISO speed: noise values are much lower and subject to more statistical variation.

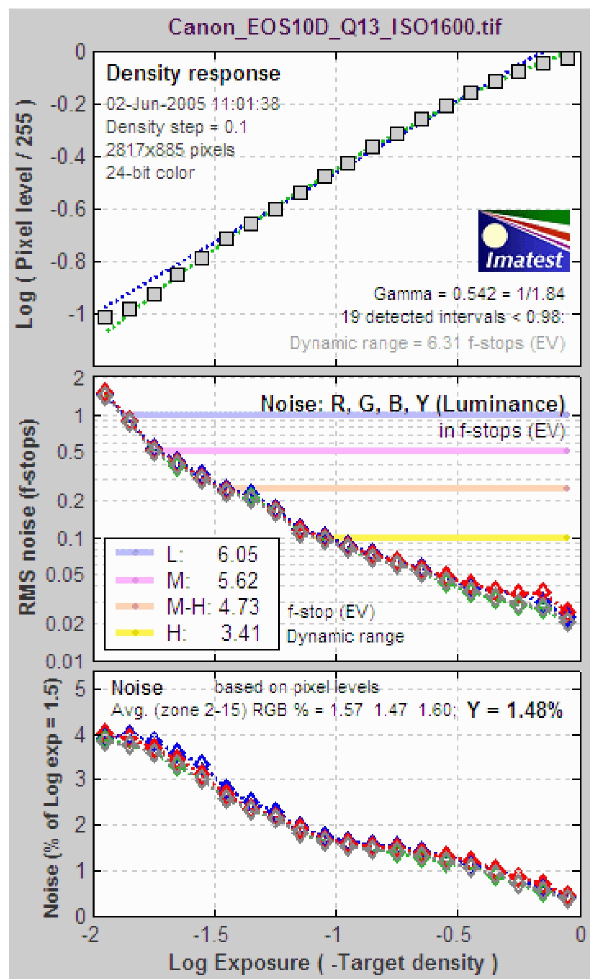

This is SNR in dB for the Canon EOS-10D at ISO 1600. (This display option was introduced with Imatest 2.3.16). SNR improves by about 6 dB for each doubling of exposure (0.301 density units); roughly 20 dB per decade (1 density unit), which is what would be expected for constant sensor noise. (This curve would be dramatically different at lower ISO speeds.)

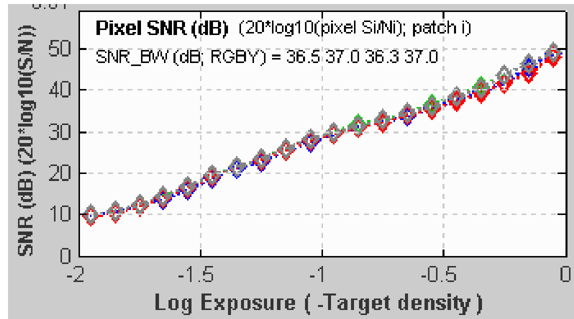

Noise summary

There are two basic types of noise.

- **Temporal.** Noise which varies randomly each time an image is captured. Measured by Stepchart and Colorcheck using two identical input images.
- **Spatial** or **fixed pattern.** Noise cause by sensor nonuniformities. Sensor designers have made heroic and largely successful efforts to minimize fixed pattern noise.

Temporal noise can be reduced by ***signal averaging***, which involves summing *N* images, then dividing by *N*. Picture Window Pro performs signal averaging using the Composite or Stack Images transformations, as described in Using Picture Window Pro in Astrophotography.

When individual images are summed *N* times, the signal pixel level increases by *N*. But since temporal noise is uncorrelated, noise ***power*** (rather than voltage or pixel level) is summed. Since voltage is proportional to the square root of power, the noise pixel level (which is proportional to noise voltage) increases by sqrt(*N*). The signal-to-noise ratio (*S*/*N* or SNR) improves by *N*/sqrt(*N*) = sqrt(N*).* When four images are averaged, *S*/*N* is improved by a factor of 2.

Several factors affect noise.

- **Pixel size.** Simply put, the larger the pixel, the more photons reach it, and hence the better the signal-to-noise ratio (SNR) for a given exposure. The number of electrons generated by the photons is proportional to the sensor area (as well as the quantum efficiency). Noise ***power*** is also proportional to the sensor area, but noise ***voltage*** is proportional to the square root of power and hence area. If you double the linear dimensions of a pixel, you double the SNR.
The electron capacity of a pixel is also proportional to its area. This directly affects dynamic range.
- **Sensor technology and manufacturing.** The biggest technology issue is CMOS vs. CCD. We won't discuss it in detail here. Until 2000 CMOS was regarded as having worse noise, but it

has improved to the point where the two technologies are comparable, differing only in detail. CMOS is less expensive because it is easy to add functionality to the sensor chip. Technology also involves other aspects of sensor design and manufacturing, all of which will improve gradually with time.

- **ISO speed.** Digital cameras control ISO speed by amplifying the signal (along with the noise) at the pixel. Hence, the higher the ISO speed the worse the noise. To fully characterize a sensor it must be tested at several ISO speeds, including the lowest and highest.
- **Exposure time.** Long exposures with dim light tend to be noisier than short exposures with bright light, i.e., reciprocity doesn't work perfectly for noise. To fully characterize a sensor it should be tested at long exposure times (several seconds, at least).
- **Digital processing.** Sensors typically have 12-bit analog-to-digital (A-to-D) converters, so digitization noise isn't usually an issue at the sensor level. But when an image is converted to an 8-bit (24-bit color) JPEG, noise increases slightly. The noise increase can be worse ("banding" can appear) if extensive image manipulation (dodging and burning) is required. Hence it is often best to convert to 16-bit (48-bit color) files. But the output file bit depth makes little difference in the measured noise of (unmanipulated) files.
- **Raw conversion.** Raw converters often apply noise reduction (lowpass filtering) and sharpening (see Noise frequency spectrum, below), whether you want it or not; even if NR and sharpening are turned off. This makes it difficult to measure the sensor's intrinsic properties.

General comments

- Imatest subtracts gradual pixel level variations from the image before calculating noise (the standard deviation of pixel levels in the region under test). This removes errors that could be caused by uneven lighting. Nevertheless, you should take care to illuminate the target as evenly as possible.
- The target used for noise measurements should be smooth and uniform– grain (in film targets) or surface roughness (in reflective targets) should not be mistaken for sensor noise. Appropriate lighting (using more than one lamp) can minimize the effects of surface roughness.

Scene-referenced noise

It is often valuable to measure noise relative top the *scene* rather than to the pixel levels. In doing so we need to take the response of the human eye into account.

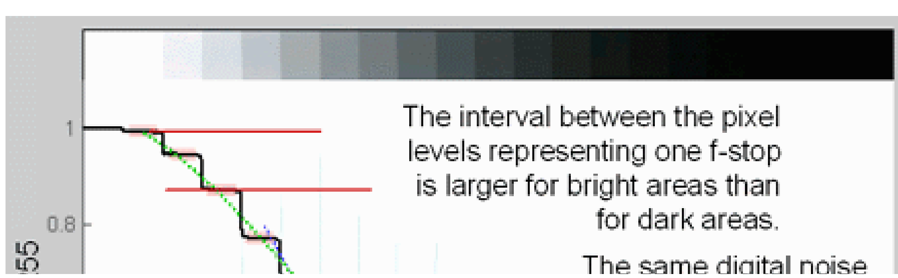

The human eye responds to *relative* luminance differences. That's why we think of exposure in terms of ***zones***, ***f-stops***, or ***EV*** (exposure value), all of which correspond to a factor of two change in exposure.

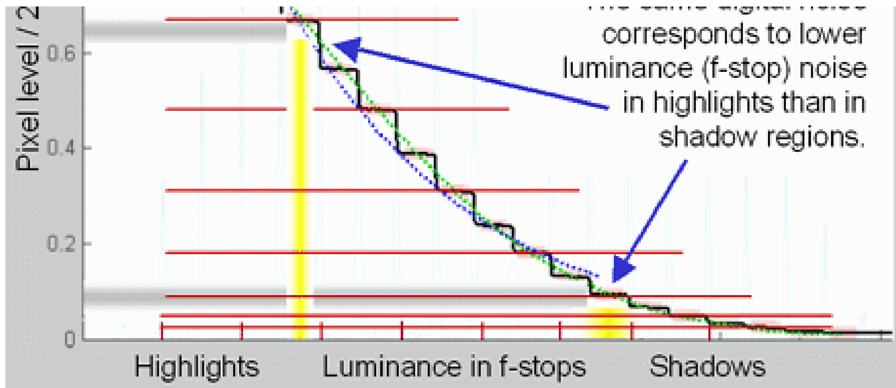

The eye's relative sensitivity is expressed by the *Weber-Fechner law*,

$\Delta L \approx 0.01\ L$ –or– $\Delta L/L \approx 0.01$

where ΔL is the smallest luminance difference the eye can distinguish. (This equation is approximate; effective ΔL tends to be larger in dark areas of scenes and prints due to visual interference from bright areas.)

F-stop noise

Expressing noise in relative luminance units, such as f-stops, corresponds more closely to the eye's response than standard pixel or voltage units. Noise in f-stops is obtained by dividing the noise in pixels by the number of pixels per f-stop. (I use "f-stop" rather than "zone" or "EV" out of habit; any of them are OK.)

Noise in f-stops = Noise in pixels / (d(pixel)/d(f-stop)) = 1/SNR_{fst}
F-stop SNR = SNR_{fst} = 1/(f-stop noise) = (d(pixel)/d(f-stop)) / Noise in pixels
where d(pixel)/d(f-stop) is the derivative of the pixel level with respect to luminance measured in f-stops ($\log_2$(luminance)). *SNR* is the Signal-to-Noise Ratio.

The above image illustrates how the pixel spacing between f-stops (and hence d(pixel)/d(f-stop)) decreases with decreasing brightness. This causes f-stop noise to increase with decreasing brightness, visible in the figures above.

Since luminance noise (measured in f-stops) is referenced to relative scene luminance, independently of electronic processing or pixel levels, it is a universal measurement that can be used to compare digital sensor quality when sensor RAW data is unavailable.

ISO 15739 noise

The updated ISO 15739 standard (due for release in 2012; the present standard was released in 2003) has several definitions for noise, SNR, and Dynamic Range that are closely related to f-stop noise. The relationships involve more than enough math to justify putting them in a green box.

Definitions:

Pxl = pixel level (same as OL = output level in the ISO standard.)
σ_{px} = noise in pixels (note: standard deviation σ is equivalent to RMS noise.)
L = Illumination or exposure level. (Units will not be important.)
f-stops = $\log_2 L$

f-stop noise = $\sigma_{fst} = \sigma_{px} /(d(Pxl)/d(\text{f-stops})) = \sigma_{px} /(d(Pxl)/d(\log_2 L))$

f-stop SNR = $SNR_{fst} = 1/\sigma_{fst} = (d(Pxl)/d(\log_2 L)) / \sigma_{px}$

We will apply the equation, $d(\log_b(x))/dx = 1/(x \ln(b))$, where $\ln(2) = 0.6931 = 1/1.4427$.

Now, from the ISO standard,

Incremental gain $= g_I = d(Pxl)/dL$ (Note linear units)
$= d(Pxl)/d(\log_2 L) \cdot d(\log_2 L) / dL = 1.442(d(Pxl)/d(\log_2 L)) / L$

Appendix D of the ISO 15739 standard defines total Signal-to-Noise Ratio as

$SNR_{ISO} = L\, g_I / \sigma_{px} = 1.4427\, L\, (d(Pxl)/d(\log_2 L)) / (L\, \sigma_{px}) = 1.4427\, (d(Pxl)/d(\log_2 L))/\sigma_{px}$

which leads to

$\boldsymbol{SNR_{ISO} = 1.4427\, SNR_{fst}}$

SNR_{ISO} is better than SNR_{fst} by a factor of 1.4427, or equivalently, 3.18 dB.

The mathematics of noise (just a taste)

Amplitude distribution

In most cases, the pixel or density variations that comprise noise can be modeled by the *normal distribution*. This is the familiar *Gaussian* or “bell” curve (**blue** on the right) whose probability density function is

$$f(x) = \exp(-(x-a)^2/2\sigma^2) / \text{sqrt}(2\pi\sigma^2)$$

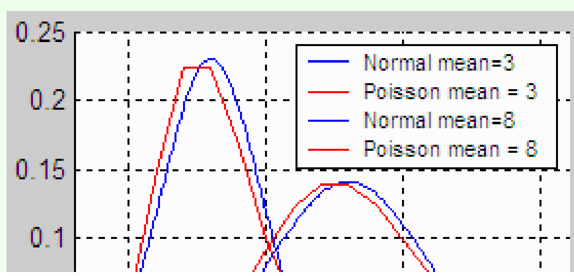

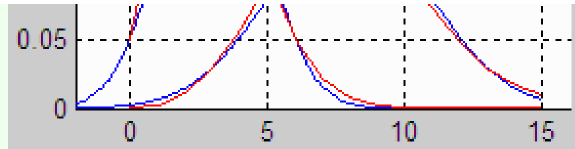

where exp(…) represents e = 2.71828… raised to a power, a is the ***mean*** (average value) of x, and σ is the ***standard deviation***. σ is proportional to the width of the distribution: for the normal distribution, about 68% of samples are between $a \pm \sigma$; 95.5% between $a \pm 2\sigma$; 99.7% between $a \pm 3\sigma$, etc.

For a set of N samples x_i with mean x_m, the *standard deviation* is

$$\sigma = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(x_i - x_m\right)^2}$$

Noise is usually measured as RMS (root mean square) voltage or pixel level, where ***RMS is equivalent to standard deviation***. RMS noise voltage is the square root of noise ***power***. The value inside the square root is also called the ***variance***.

Noise varies with the pixel level, i.e., it differs in areas with different tonal values.

The normal curve arises from a remarkable mathematical result called the *central limit theorem*: When a variable (such as voltage or pixel level) is affected by a large number of perturbations, the overall density function approaches the normal curve, regardless of the distributions of the individual perturbations. This is why the normal curve is by far the most common probability distrubution.

But the normal distribution doesn't apply in all situations. For low light levels (low photon counts), where the normal distribution could result in negative counts, the *Poisson distribution* (**red** in the illustration above) gives the correct result.

$f(s) = \exp(-m)m^s/s!$

where m is the mean, $s \geq 0$ is an integer, and $s!$ = "s factorial" = $(s)(s\text{-}1)(s\text{-}2)\ldots(1)$. The standard deviation is $\sigma = \text{sqrt}(m)$. Shot (photon) noise, described in the Kodak publication, CCD Image Sensor Noise Sources, has Poisson statistics. For large values of m, the Poisson distribution approaches a normal distribution.

Noise frequency spectrum

In addition to an amplitude distribution, noise is characterized by a ***frequency spectrum***, calculated by taking the ***Fourier transform*** of the spatial image. The spectrum is closely related to appearance.

Two spectra are shown below, taken from the second Stepchart figure. The first is for the image from column (B), above, which contains simulated white noise. The second is for the image in column (D), above, taken with the Canon EOS-10D at ISO 1600. The images shown below have been magnified 2X (using the nearest neighbor resizing algorithm) to emphasize the pixel distribution. They are close approximations to the images used to calculate the spectra (but ***not*** the exact images).

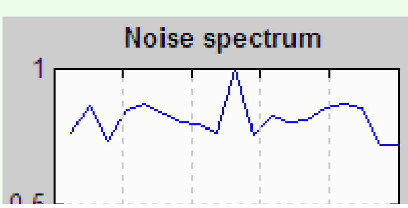

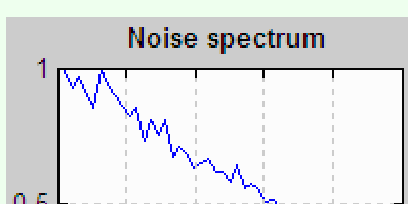

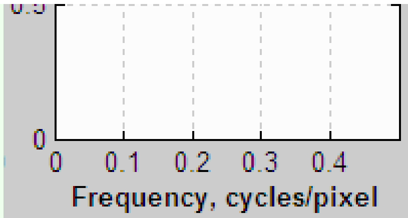

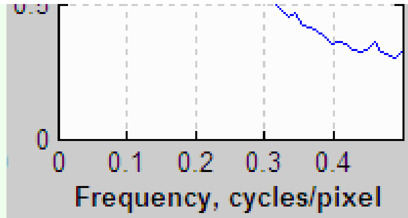

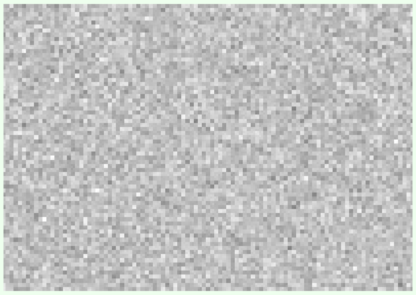

White noise, enlarged 2X (nearest neighbor)

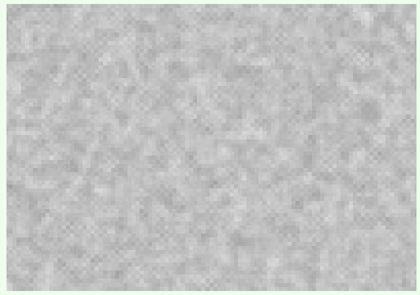

Blurred noise, enlarged 2X (nearest neighbor)

White noise has two key properties. 1. The values of neighboring pixels are *uncorrelated*, i.e., independent of one-another. 2. Its spectrum is *flat*. The white noise spectrum (above, left) shows statistical variation and a small peak at 0.25 cycles/pixel, probably caused when the image in column (B) was resized for display.

For spectral (non-white) noise, neighboring pixels are correlated and the spectrum is not flat. The spectrum and image (above, right) are the result of blurring (also called smoothing or lowpass filtering), which can result from two causes. 1. The Bayer sensor demosaicing algorithm in the RAW converter causes the noise spectrum to drop by about half at the Nyquist frequency (0.5 cycles/pixel), and 2. Noise reduction (NR) software lowpass filters noise, i.e., reduces high frequency components. NR usually operates with a threshold that prevents portions of the image near contrast boundaries from blurring. But NR comes at a price: detail with low contrast and high spatial frequencies can be lost. This causes the “plasticy” appearance sometimes visible on skin. Some people love it; I don’t. (Plastic surgeons make a ***lot*** more income than I do.)

The visibility of noise depends on the noise spectrum, though the exact relationship is complex. Noise at high spatial frequencies may be invisible in small prints (low magnifications) but very damaging in large prints (large magnifications). Because of the complex nature of the relationship, Kodak has established a ***subjective*** measurement of grain (i.e., noise) called Print Grain Index (Kodak Technical Publication E-58).

Sharpening and unsharp masking (USM) are the inverse of blurring. They boost portions of the spectrum and cause neighboring pixels to become negatively correlated, i.e., they exaggerate the differences between pixels, making the image look noisier. Unsharp masking is often applied with a threshold that restricts sharpening to the vicinity of contrast boundaries. This prevents noise from degrading the appearance of smooth areas like skies.

When poor quality lenses are used (or the image is misfocused or shaken), the image is lowpass filtered (blurred) but the noise is not. Some sharpness loss can be recovered with sharpening or

USM, but noise is boosted in the process. That’s why good lenses are important, even when digital sharpening is available.

RAW converters often perform both sharpening and NR, whether you want it or not. This makes it difficult to compare the intrinsic performance of different cameras. Standardized sharpening is an imperfect attempt compensate for these software differences.

Links

CCD Image Sensor Noise Sources (a Truesense (formerly Kodak) Reference Document) An excellent, highly readable introduction to the different noise sources. Noise sources in CMOS sensors differ only in relative magnitude.

Neat Image is a remarkable program for reducing noise. Instead of lowpass filtering, the software is trained to recognize the structure of the noise. It can reduce noise with minimal side-effects. Particularly valuable at high ISO speeds.

EMVA 1288 Standard contains an excellent, well-written introduction to the mathematics of noise.

Photography — Electronic still-picture imaging — Noise measurements (ISO 15739:2003) Murky, difficult to read, and expensive (106 Swiss Francs), but it’s the official standard. Reads like it was written by a committee because it was (ISO/TC 42). An update is due in 2012. Frequently refers to other ISO standards. Its Dynamic Range measurement is extrapolated from measurements from a black reference density of 2.0.