A Closer Look at Dynamic Range and Signal to Noise Ratio in Spectrometers

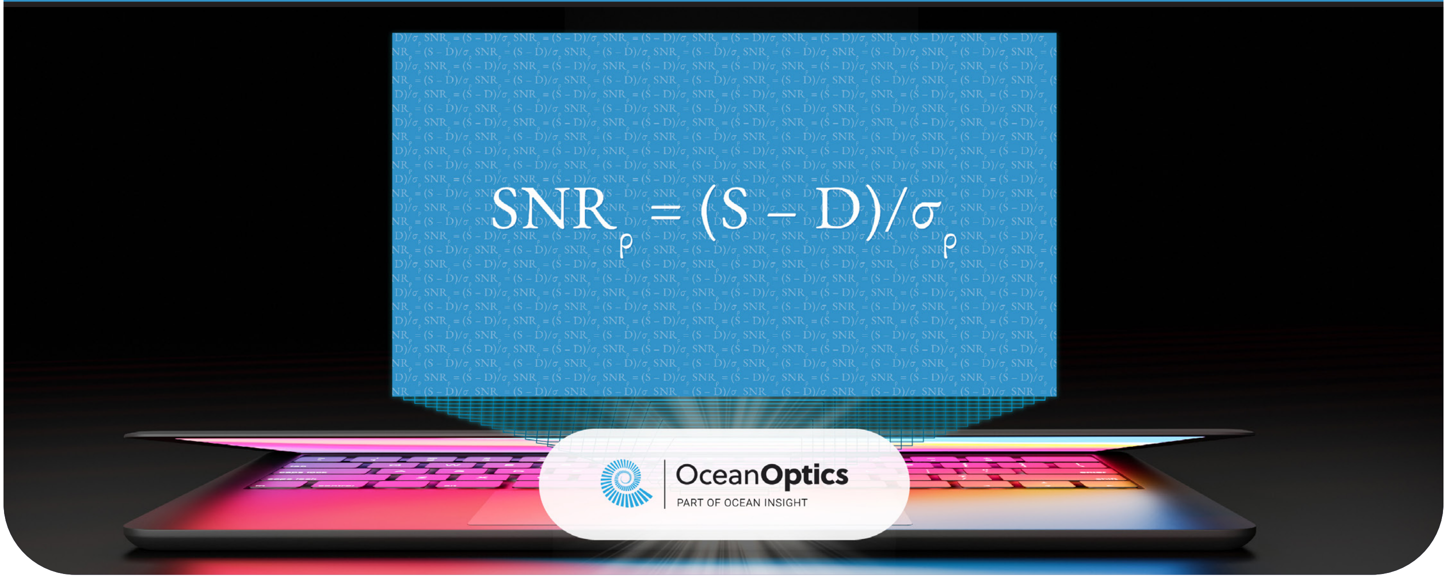

Spectrometer performance is characterized by benchmarks including spectral range, optical resolution and stray light. In this tech tip, we consider the impact of the performance criteria dynamic range and signal to noise ratio.

Spectroscopy is a complex technique, with many variations and nuances to consider, often framed in terminology and specifications that are not familiar to users or can be interpreted in different ways. Also, users often have very distinct measurement needs – for example, as different as characterizing a laser beam or measuring the fluorescence of biological dyes – which determine the relevance of each criteria.

Within that context, we will focus in this technical tip on practical definitions of dynamic range and signal to noise ratio (SNR), which are common spectrometer specifications, and weigh the importance of each in evaluating spectrometer performance.

Dynamic Range

In spectroscopy, dynamic range is the ratio between the maximum and minimum signal intensities that a spectrometer can detect. More specifically, dynamic range is the maximum detectable signal (i.e., near saturation) divided by the minimum detectable signal. The minimum detectable signal is defined as the signal with an average equal to the baseline noise (i.e., the noise from detector readout, electronics and so on).

Why is dynamic range relevant? For spectroscopy, dynamic range gives you an idea of how much room on the scale is available to detect spectral features in a single scan. More simply, a high dynamic range allows you to see both "small" peaks (i.e., lower on the y-axis, near the baseline) and tall peaks (i.e., toward the top of the y-axis) clearly within the same measurement. With lower dynamic range, the small peaks would be impossible to discern because of the noise around the baseline.

If you think of dynamic range in terms of sound, an average healthy person with good hearing has an auditory system with dynamic range sufficient to hear something as soft as a whisper and as loud as a jet engine.

For Ocean Optics spectrometers, we report dynamic range in terms of a single acquisition, which is defined as the shortest integration time giving the highest possible dynamic range. The dynamic range specification of the entire system is defined as the ratio of maximum to minimum signal at the longest integration time and the ratio of the maximum to minimum integration time:

$DR_{\text{single acquisition}}$ = number of counts at saturation/baseline noise at shortest integration time

DR_{sys} = (number of counts at saturation/baseline noise at longest integration) x (longest integration time/shortest integration time)

Here's where things get a little tricky. Although dynamic range is typically reported for a single scan, you can always increase dynamic range by averaging over multiple scans. But that is not especially practical, as dynamic range does not include noise factors related to detector well depth (dark current) and shot noise (fluctuations in the numbers of photons detected).

In our experience, users sometimes conflate dynamic range and spectrometer sensitivity, which is the minimum amount of light that can be detected at a fixed integration time for a given spectrometer configuration. Although dynamic range and sensitivity are related, they are not interchangeable terms. Indeed, sensitivity is affected by many factors and is a topic for a separate technical paper.

Signal to Noise Ratio (SNR)

SNR and dynamic range cover similar territory but are different parameters. SNR is typically defined as the maximum signal intensity divided by the noise intensity at a particular signal level, which means it can change from measurement to measurement. Here's another way of thinking of it: If dynamic range is the maximum signal divided by the noise, SNR is the maximum signal divided by the noise at saturation.

Some folks may not realize that noise increases as a function of signal; i.e., the signal itself begins to add noise to the system. Since this noise increases as a function of signal due to photon noise, the SNR function is a plot of individual SNR values versus the signal at which they were obtained. Typically, manufacturers report the SNR specification as the maximum signal divided by the noise at the minimum integration time for a single scan (more on this later).

Often, SNR testing is conducted with a broadband light source so that the spectral peak nearly saturates at the lowest detector integration time or at an integration time well below the thermal noise limit (the spectrum should also have an area of low or nearly zero counts). SNR is calculated by taking X scans without light (we use 100 scans) and calculating the mean baseline count value at each pixel, then taking X scans with light and calculating the mean and standard deviation of each pixel output count.

$$SNR\rho = (S - D)/\sigma\rho$$

S = mean intensity of the samples (with light)
D = mean of dark (without light)
σ = standard deviation of samples (with light)
ρ = pixel number

The complete SNR versus signal graph is a plot of the calculated SNRρ values (the noise) versus Sρ – Dρ (the signal). This covers a wide range of peak counts (from dark to nearly saturated). Since all of the pixels have the same response curve, the data for the SNR versus signal graph can come from all of the different pixels.

So, how relevant is SNR for my application? SNR is one way to determine what the uncertainty of an individual measurement will be. For example, imagine you're making an absolute irradiance measurement and want to know the brightness (intensity) of an LED spectral peak. If you're SNR is 100:1, that means you have 1% uncertainty in measuring that peak (the signal is 100x the noise).

Other examples of applications where SNR matters include absorbance concentration measurements, especially when you want to detect a wide range of concentration levels of the analyte of interest; and applications that require high accuracy, particularly when the level of incident light on the detector is low as a result of low reflectivity, low transmission or other factors that limit incident light or integration time.

SNR and Signal Averaging

Users can take steps to improve SNR. They can reduce noise in their setup and/or add more signal by a) increasing light source output, b) increasing the integration time of the spectrometer, or c) limiting the incoming lamp spectrum to only the wavelength area of interest. The latter step can be helpful because it takes advantage of the full dynamic range of the detector at the wavelengths where it matters most, as the signal has less intensity at the edges of the spectrum.

There is a caveat to these tactics. For several reasons, some or all of these steps may not be practical or may introduce unwelcome tradeoffs. Instead, users can improve SNR by using signal averaging. For example, for time-based averaging performed in software, the SNR will increase by the square root of the number of spectral scans used. SNR of 300:1, for example, will become 3000:1 if 100 scans are averaged. For spatially based averaging (boxcar averaging), the SNR will increase by the square root of the number of detector pixels averaged.

Signal averaging methods are effective in increasing SNR, but because most manufacturers limit the definition of SNR to a single scan, the specification lacks context. If manufacturers report SNR as the maximum signal divided by the noise at the minimum integration time for a single scan, what is that minimum integration time? And how much light was needed to achieve that particular SNR?

For example, if you have a highly sensitive spectrometer that can get all the light it needs in a millisecond, why not measure 100 scans versus a single scan? That would increase the SNR by 10x and takes only 100 milliseconds. Unless you're measuring a transient event, that integration time would suffice for most applications.

There's another issue to consider: Not all spectrometer manufacturers report SNR using the same criteria. Historically, Ocean Optics has specified SNR without applying signal averaging to boost the ratio, while some competitors have used signal averaging to artificially inflate SNR values on inferior spectrometers.

For these reasons, Ocean Optics has begun to report SNR as the maximum SNR per second. This is a more meaningful parameter because it provides a more definitive unit of measure (i.e., one second versus one scan).

SNR and High Speed Averaging Mode

Until recently, the only signal averaging method available to improve SNR with Ocean Optics spectrometers required averaging the spectra in OceanView software on the host computer. This limits the number of scans you might be able to get out of the device each second because of delays in the microprocessor and other computing overhead.

High Speed Averaging Mode (HSAM) is a new function in Ocean Optics spectrometers that offers users a novel method to dramatically enhance SNR. HSAM is a hardware-accelerated averaging capability available in Ocean ST, SR and HR spectrometers, and to operate requires OceanDirect, a device driver platform with API.

HSAM enables you to perform signal averaging on the spectrometer itself. So, because we're not exporting the data for each cycle (scan), we can collect significantly more data over the same period compared with averaging via software. This results in a much higher SNR per unit time versus traditional averaging.

One way to consider signal averaging is this: If you're using signal averaging and want it to happen faster, use HSAM. Here are two examples where HSAM can add value:

First, imagine you're making color measurements of flat samples on a conveyor belt, as part of a QC process. The samples are illuminated with a very bright light, with the goal of making a reflected color measurement on the samples at a certain accuracy, as rapidly as possible. The faster you can make your measurements, the better run rate you can achieve, because that's the bottleneck in your process.

Second, imagine you're measuring subtle absorbance changes in a sample that's mostly transparent and absorbs very little. You have high signal and are looking for small changes in absorbance, so SNR is a limiting factor. Signal averaging can help detect those small absorbance peaks, and HSAM can get the data out faster that enables that detection.

There is a tradeoff with HSAM: You won't be able to observe each spectra. This is relevant in kinetic measurements, for example, where your goal is to monitor changes that occur very rapidly. With HSAM, you would miss that data, because you've taken 100 scans, for example, and averaged them to a single scan. This is different than our earlier absorbance measurement example, where we needed very high SNR and needed the data fast.

Summary

Ultimately, dynamic range and SNR provide important performance criteria in consideration of your spectrometer and spectral configuration choices. But they should be viewed within the larger context of what you're trying to achieve with your measurements and why it's important not to overestimate the impact of any one specification. Your setup really is no greater than the sum of its parts.