

How to read an FFT spectrum: resolution, windows, peaks

After reading, you can compute frequency resolution for a vibration record, justify your window choice, and read spectral peaks as clues rather than diagnoses, using a worked example.

In short

- Frequency resolution is $\Delta f = f_s / N$, so a longer record gives finer spectral lines and better separation of close peaks.
- A window function reduces leakage but widens peaks, so the choice should follow the goal of the measurement.
- Spectral peaks are clues about machine behaviour; a diagnosis needs more measurements and the operating context.

01 / Time domain and frequency domain

An accelerometer or velocity sensor records a vibration signal that changes over time. This record lives in the time domain: each sample is the amplitude at one instant. Faults such as gear mesh, bearing defects or unbalance, however, usually put energy at specific frequencies, and that information is easy to miss in a waveform plot. This distinction is the basis for knowing which information a record carries.

The fast Fourier transform (FFT) converts an N-sample record into the frequency domain. The result is a spectrum: the amplitude of each frequency line. From N samples you get roughly N/2 usable lines. The highest valid frequency is half the sampling rate, $f_s/2$. Components above that limit fold back into the spectrum at a wrong frequency; this effect is called aliasing.

02 / Resolution: how Δf is set

Frequency resolution is the spacing between neighbouring frequency lines. It equals $\Delta f = f_s / N$, where f_s is the sampling rate and N the number of samples. Because the record duration is $T = N / f_s$, the same result can be written as $\Delta f = 1 / T$. Resolution therefore improves as the record gets longer.

The practical consequence is simple. With N fixed, raising the sampling rate makes Δf larger and the spectrum coarser. The usual way to sharpen the lines is a longer record. Lowering f_s also reduces Δf , but it lowers the $f_s/2$ limit at the same time, so the analysis band and the sampling rate should be chosen together.

03 / Window functions and spectral leakage

An FFT treats the record as if it repeats forever. If the record does not contain a whole number of periods, there is a discontinuity between its start and its end. That discontinuity spreads energy into neighbouring frequencies; this is spectral leakage. A window function tapers the edges of the record toward zero, which reduces the effect.

The Hann window is a common starting point for general vibration work. It reduces leakage noticeably but widens each peak, which can make two nearby components harder to separate. A rectangular window suits records that contain an exact whole number of periods, or a short transient such as an impact. Welch's method splits the record into overlapping segments and averages their spectra. This lowers the variance of a noisy spectrum, but the resolution is then set by the segment length rather than by the full record.

04 / Worked example: where the 1× peak falls

Hypothetical example. Suppose $f_s = 10\,240$ Hz and $N = 8\,192$ samples. Then $\Delta f = 10\,240 / 8\,192 = 1.25$ Hz. The record duration is $T = 8\,192 / 10\,240 = 0.8$ s, and $1 / 0.8 = 1.25$ Hz, which agrees.

If the shaft runs at 1 480 rpm, the 1× frequency is $1\,480 / 60 = 24.67$ Hz. Its line position is $(1\,480 / 60) / 1.25 = 19.73$. So the peak falls between line 19 (23.75 Hz) and line 20 (25.00 Hz). Because the true peak does not sit exactly on a line, its energy spreads over two lines and a small amount of leakage is expected. This is a property of the frequency grid, not a sign of a faulty signal. When you read the peak, note the shaft speed with the record and look at both neighbouring lines together.

05 / Reading peaks and their limits

The strongest peaks are often multiples of shaft speed. A 1× peak can accompany unbalance and a 2× peak can accompany misalignment, though either can appear for other reasons. Harmonics show up as 3×, 4× and higher multiples. Sidebands, which are pairs of lines around a peak, can indicate periodic modulation. All of these are clues, not diagnoses. The same peak can come from looseness, a bearing defect, a gear problem or an electrical source.

If the machine changes speed during the record, for example during run-up, a single FFT does not show when each peak occurred. A spectrogram, which shows the spectrum over successive time slices, makes that change visible. A short record with little averaging also gives a noisy spectrum, and small peaks may not be separable from the noise.

06 / How to start with your own data

Write the measurement context beside every record: sampling rate, sample count, window type, shaft speed, measurement point and load. Without this information you cannot later confirm whether a peak is 1× or some other component.

A practical order of work:

- Choose the analysis band first, then set f_s for that band.
- Compute Δf and check whether the expected peaks are resolved.
- Record several times at the same load and speed, and check whether the peaks repeat.
- Compare the result with an earlier record from the same point, taken while the machine was known to be in good condition.

Keep the raw record unchanged, so the same data can be analysed again if the window, the band or the averaging changes. At this stage, measurement discipline usually decides the quality of the result more than the choice of algorithm does.

Try it

Open one of the built-in synthetic signals in spectra, change the window setting, and compare how the width of the strongest peak and its side leakage change.

References and online guide

[Open the online guide](#)

[Fast Fourier transform \(Wikipedia\)](#)

[Window function \(Wikipedia\)](#)

[Spectral leakage \(Wikipedia\)](#)

[welch \(SciPy\)](#)

Proposed pilot steps do not claim field accuracy. Synthetic examples are not measurements from a production site.