Frequency Response and Analysis Techniques

FRF · measurement · diagnostics · standards

Abu Haqeem

MEC521 Vibrations · Handbook Chapter 4 · Weeks 12–14

Semester 20264

How to use this chapter

Three documents, three different jobs. Use the right one.

These slides	Where every formula *comes from*, and what it means. Follow the derivations here.
The handbook	*Nota padat* — the compact reference. Look up a result, a tag, a worked example. Not a first read.
The tutorial	Where you find out whether you actually understood it.

This chapter is where the course meets the instrument

Chapters 1–3 predicted what a machine *should* do. Chapter 4 is about measuring what it *does*, and deciding what to do about it.

Where Chapter 4 is going

1. **From zero** — what a frequency, a spectrum and the FFT actually are. start here
2. **The frequency response function** — the measured cousin of everything in Chapter 2. §4.1
3. **Getting a spectrum right** — sampling, aliasing, leakage, averaging — and the **hammer test**, from a tap to a list of modes. §4.2
4. **Measuring properly** — sensors, mounting, and which quantity to record. §4.3
5. **Reading a spectrum for faults** — shaft orders, bearings, envelopes. §4.4
6. **Deciding** — severity standards and what to recommend. §4.5

This is the chapter your Mini Project is built on.

Running the argument backwards

Chapters 1–3

Given m, c, k — predict the response.

Design.

Chapter 4

Given the measured response — infer m, c, k, and whether anything is wrong.

Diagnosis.

Same physics, opposite direction

Every resonance you learned to calculate is a peak someone can measure. Every peak someone measures is a resonance you can now explain.

And the inverse problem is the harder one — measurements are noisy, incomplete, and occasionally lie.

You already do frequency analysis

With your ears

- you hear a guitar chord and can tell it is *several* notes;
- a mechanic hears "a whine" and "a knock" in one engine and knows they have different causes;
- the bars jumping on a music player are a spectrum, live.

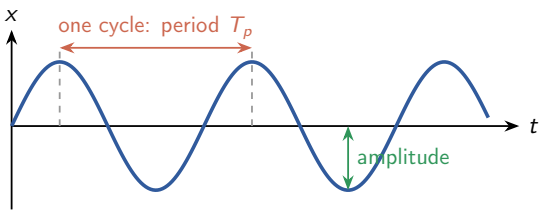
With machines

- a ceiling fan at 200 rev/min turns 3.33 times a second; its 3 blades pass the light 10 times a second;
- a car engine at 3000 rev/min: 50 times a second;
- a transformer hums at 100 Hz: twice the 50 Hz TNB supply.

The engineering question

A sensor on a machine records one wiggly line. Hidden inside it are many rhythms, each with its own cause. Chapter 4 is about pulling them apart — and that is all a *spectrum* is.

Frequency, in one picture



Frequency is how many cycles fit in one second:

$$f = \frac{1}{T_p} \quad [\text{Hz} = \text{cycles per second}]$$

A shaft speed in rev/min becomes hertz by dividing by 60:

$$3000\,\text{rev/min} \;\rightarrow\; \frac{3000}{60} = 50\,\text{Hz}$$

Chapters 1–3 used $\omega = 2\pi f$ in $\text{rad}\,\text{s}^{-1}$. An analyser always shows f in Hz. Same thing, different unit.

Every real signal is a mixture

A real vibration is a sum of simple rhythms, each with a size:

$$x(t) = A_1 \sin 2\pi f_1 t + A_2 \sin 2\pi f_2 t + A_3 \sin 2\pi f_3 t + \dots$$

Time view

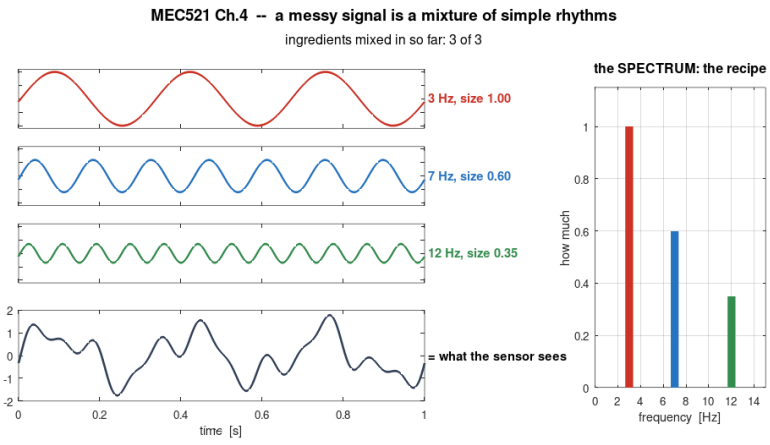
the wiggle the sensor records. Everything is in it, but tangled — like a cake.

Frequency view (the spectrum)

one bar per ingredient: *which* frequency, *how much*. The recipe.

Same information, two views. The spectrum is easier to read, because each bar points at one cause: a shaft, a gear, a blade, a bearing.

Watch a recipe being mixed



Three ingredients (3, 7 and 12 Hz) go in one at a time; the sensor only ever sees the bottom line. GIF:

`fft-anim/MEC521-ch4-chord.gif`.

How the FFT finds the ingredients

The analyser only has the tangled line. To ask *"is there any* 7 Hz *in here?"* it does three things:

1. make a clean **test rhythm** at 7 Hz;
2. **multiply** it by the signal, point by point;
3. **average** the product.

If 7 Hz is in there

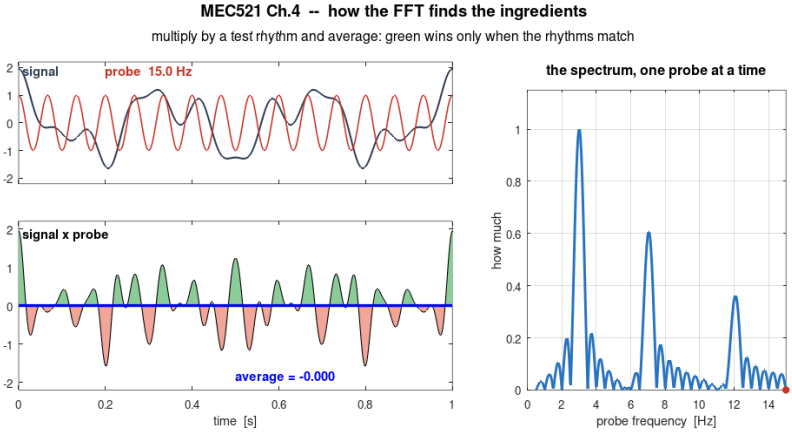
the two rise and fall together: the product is mostly positive, and the average is large — it equals the ingredient's size.

If it is not

the product is positive as often as negative, and the average cancels to zero.

Repeat for every frequency and you have the spectrum. The **FFT** (fast Fourier transform, Cooley and Tukey, 1965) is just a very fast way of doing all those tests at once. It uses a sine *and* a cosine test rhythm, so it does not matter when each ingredient started.

Watch the test rhythm sweep



Green area above the line, red below. They only stop cancelling when the test rhythm matches an ingredient — and the spectrum on the right grows a peak exactly there. GIF: `fft-anim/MEC521-ch4-probe.gif`.

The computer only sees dots

An analyser, a phone or a data logger does not record a smooth line. It **samples**: it writes down the value f_s times a second.

$$N \text{ samples at } f_s \quad \Longrightarrow \quad \text{a record } T = \frac{N}{f_s} \text{ seconds long}$$

Two rules follow from this, and every spectrum you will ever read obeys them:

rule	**in words**	**in § 4.2**
$\Delta f = 1/T$	to tell close rhythms apart, *listen longer*	resolution
$f_{\max} = f_s/2$	you need at least two dots per cycle	aliasing

Try it tonight: the free *phyphox* app shows your phone accelerometer's spectrum live. Put the phone on a running washing machine.

The frequency response function

$$H(\omega) = \frac{X(\omega)}{F(\omega)} \qquad \text{F4.1.1}$$

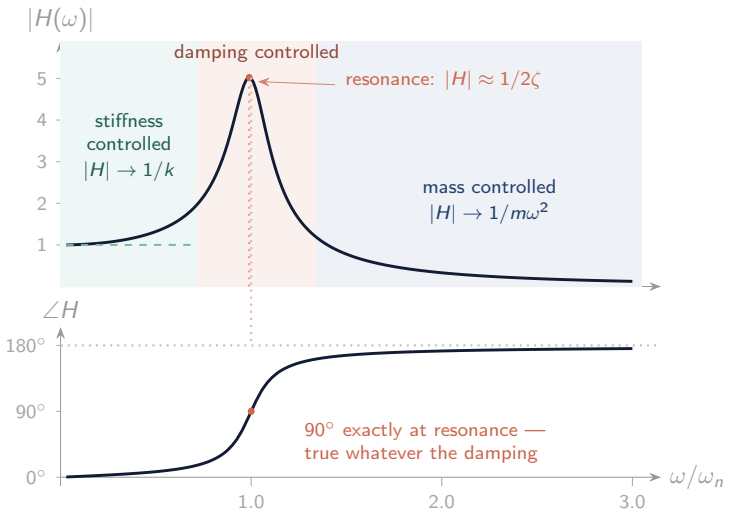
One complex number at each frequency: its **magnitude** is the amplitude ratio, its **argument** is the phase lag.

You have already met it

The magnification factor of Chapter 2 *is* $|H|$, scaled by k. The phase curve is $\arg H$. Chapter 2 derived it; Chapter 4 measures it.

Measured by hitting the structure with an instrumented hammer and recording both the force and the response.

Anatomy of an FRF



Three regions, and what each one tells you

$$|H| \to \frac{1}{k}\ (\omega \to 0), \qquad |H|_{\text{peak}} \approx \frac{1}{c\omega_n}, \qquad |H| \to \frac{1}{m\omega^2}\ (\omega \to \infty)$$

F4.1.2

Region	Asymptote	What you can read off
low frequency	$1/k$	the **stiffness**
at the peak	$1/c\omega_n$	the **damping**
high frequency	$1/m\omega^2$	the **mass**

One measured curve contains all three system parameters.

Three names for the same idea

Which quantity you divide by the force decides the name — and the units.

Name	Definition	Units
receptance	X/F	$\mathrm{m\,N^{-1}}$
mobility	$\mathrm{j}\omega H$	$\mathrm{m\,s^{-1}\,N^{-1}}$
accelerance / inertance	$-\omega^2 H$	$\mathrm{m\,s^{-2}\,N^{-1}}$

F4.1.3

They are the same measurement

Each is the previous one multiplied by $\mathrm{j}\omega$. Instruments report whichever the analyser was set to — always check which before comparing two plots.

The FRF of a two-degree-of-freedom system

Two natural frequencies means **two peaks**, and between them an *anti-resonance* — a frequency at which that particular point barely moves at all.

Peaks are a property of the system; anti-resonances are not

Move the accelerometer and the peaks stay where they are. The anti-resonances move, because they depend on where you measure and where you excite.

Which is why a single measurement point can miss a mode

If the sensor sits at a node of a mode, that mode simply does not appear. Two sensors, or two hit points, are the cheapest insurance in modal testing.

From a time record to a spectrum

The discrete Fourier transform takes N samples of $x(t)$ and returns the amplitude at $N/2$ discrete frequencies.

Three numbers govern everything

$$T = \frac{N}{f_s}, \qquad \Delta f = \frac{f_s}{N} = \frac{1}{T}, \qquad f_{\mathrm{Nyq}} = \frac{f_s}{2}$$

F4.2.1 F4.2.2

$\Delta f = 1/T$ is the one to remember: **resolution is set by how long you record**, not by how fast you sample.

Three acquisition numbers — and the trade they force

Want	Change	Cost
higher top frequency	raise f_s	coarser Δf at fixed N
finer resolution	record longer	the machine must stay steady
both	raise N	more memory, more time

Resolution is not free

To separate two lines 1.75 Hz apart you need Δf below about 0.875 Hz, hence $T > 1.1$ s of *steady* running. On a machine that is changing speed, that record does not exist.

Why resolution needs time: listen for the beat

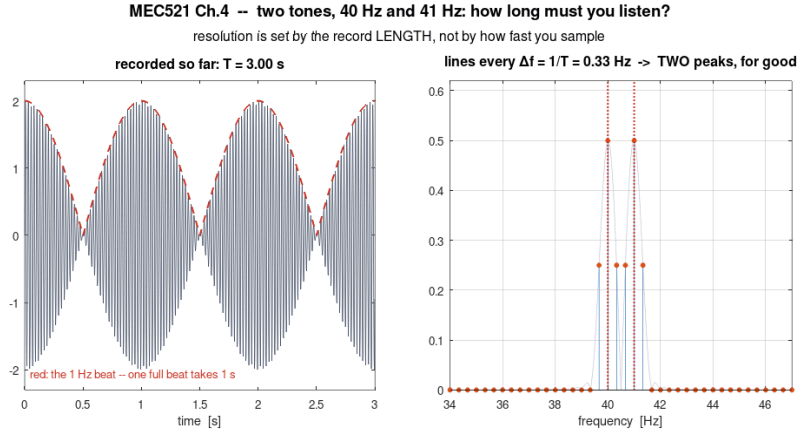
Play 40 Hz and 41 Hz together and the sum swells and fades once a second — a **beat**. Until you have heard at least one beat, nothing tells you there are two tones and not one.

That is $\Delta f = 1/T$ in words: a 1 s record gives lines 1 Hz apart, just enough to *start* telling them apart.

In practice, about twice that

With the Hann window an analyser uses, the two peaks separate for good only once $T \geq 2.30\,\mathrm{s}$ — just over twice the beat period. The rule on the previous slide (make Δf half the spacing) is that factor of two.

Watch the two peaks come apart



Near $T = 1.2\,\mathrm{s}$ a dip appears and vanishes again: two tones that start together can look separate by luck of timing. Trust the spectrum only once they stay apart. GIF: `fft-anim/MEC521-ch4-res.gif`.

Worked example — acquisition settings

A gearbox is monitored at $f_s = 5120\,\text{Hz}$ with $N = 8192$ samples.

Step 1 — record length

$$T = \frac{N}{f_s} = \frac{8192}{5120} = 1.6\,\text{s}$$

Step 2 — resolution

$$\Delta f = \frac{f_s}{N} = \frac{5120}{8192} = 0.625\,\text{Hz} \qquad \left(= \frac{1}{T}\ \checkmark\right)$$

Step 3 — the highest frequency that is real

$$f_{\text{Nyq}} = \frac{5120}{2} = 2560\,\text{Hz}$$

Step 4 — what happens to a genuine 3000 Hz **component?** It is above Nyquist. It does not disappear — it **aliases** to

$$|3000 - 5120| = 2120\,\text{Hz}$$

and appears there as a line that looks entirely real.

Aliasing: the one error you cannot undo

$$f_{\text{alias}} = |f - nf_s| \qquad \text{F4.2.3}$$

Sampling too slowly does not lose the high-frequency content. It **folds it down** into the band you are looking at, where it is indistinguishable from a real component.

No amount of post-processing recovers the truth

Leakage can be reduced afterwards by windowing. Noise can be reduced by averaging. **Aliasing cannot be removed at all** — the information is gone.

Which is why every serious analyser has an **analogue** anti-aliasing filter ahead of the converter — it must act before sampling, not after.

You have seen aliasing on your phone

A phone video is a sampler: 24 to 30 pictures a second. Film anything that spins faster than half that rate and the video lies.

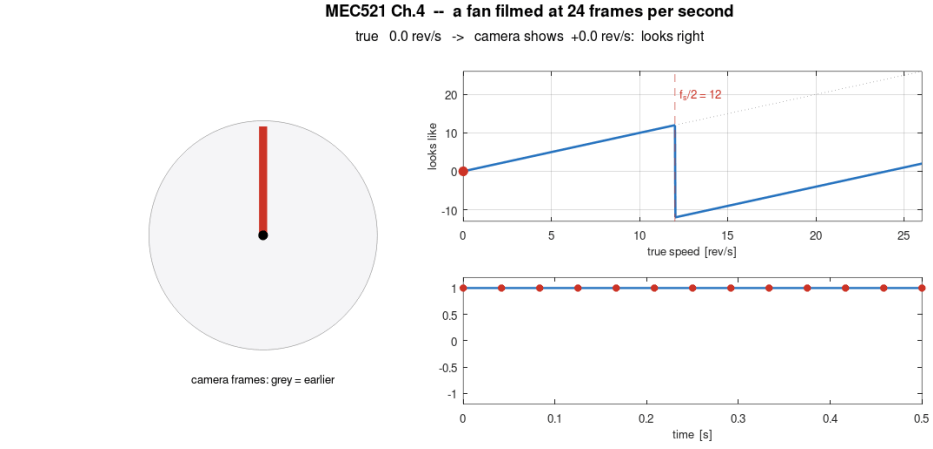
- car wheels that seem to roll **backwards** in films;
- helicopter rotors that look **frozen** in online clips;
- a fan that seems to turn **slowly** although it is at full speed.

The fan at 24 frames per second ($f_s/2 = 12\,\text{rev/s}$)

true speed	the video shows	by $f_{\text{alias}} = \lvert f - n f_s \rvert$
$10\,\text{rev/s}$	$10\,\text{rev/s}$	below $f_s/2$: honest
$18\,\text{rev/s}$	$6\,\text{rev/s}$, **backwards**	$\lvert 18 - 24 \rvert = 6$
$24\,\text{rev/s}$	standing still	$\lvert 24 - 24 \rvert = 0$

The pictures of 18 and $-6\,\text{rev/s}$ are identical, frame for frame. No software can tell which was real — that is why the filter must come *before* the sampler.

Watch the fan lie



Bottom: the red dots are what the camera keeps; the blue wave is the slower motion they also fit perfectly. GIF:

`fft-anim/MEC521-ch4-alias.gif`.

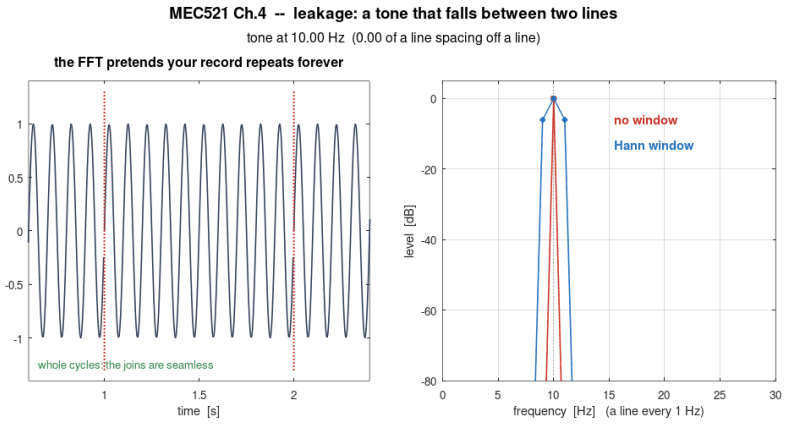
Leakage, and why windows exist

The DFT assumes your record repeats exactly. If a sinusoid does not fit a whole number of cycles into T, the assumed repetition has a step in it — and a step contains every frequency. The energy of one clean line **leaks** into its neighbours: the peak reads low and the skirts read high.

Window	**Use for**	**Trade**
Rectangular (none)	transients, hammer impacts	worst leakage
Hanning	general continuous signals	the default choice
Flat-top	accurate *amplitude* readings	poor frequency resolution

A window trades amplitude accuracy against frequency accuracy. You cannot have both.

Watch a tone leak



A 10.5 Hz tone in a 1 s record (lines every 1 Hz): with no window the peak reads 3.7 dB low and ten lines away it is still at −32.6 dB; the Hann window reads 1.4 dB low and −68.8 dB. GIF: `fft-anim/MEC521-ch4-leak.gif`.

Averaging: what it fixes and what it cannot

$$\text{noise floor} \propto \frac{1}{\sqrt{n_{\text{avg}}}} \qquad \text{F4.2.4}$$

Averaging n spectra drops the *random* noise floor by $\sqrt{n}$, so buried periodic lines rise out of it.

n_{avg}	4	16	64
noise floor	$\div 2$	$\div 4$	$\div 8$

Diminishing returns, and a hidden cost

Going from 16 to 64 averages costs four times the record time for a factor of two. And the machine must stay at *constant speed* throughout, or the lines smear.

Tapping to listen is older than any analyser

- you knock a **watermelon** at the pasar to judge it;
- a contractor taps **floor tiles** to find the hollow ones;
- railway **wheel tappers** once struck every train wheel with a hammer — a cracked wheel rings dull.

In each case a short knock shakes the object at *all* frequencies at once, and the object answers only at the frequencies it likes: its **natural frequencies**.

The engineering question

Which frequencies does this structure ring at, how quickly does each one die away, and in what shape does it move? A **hammer test** answers all three, with a measured force instead of a guess.

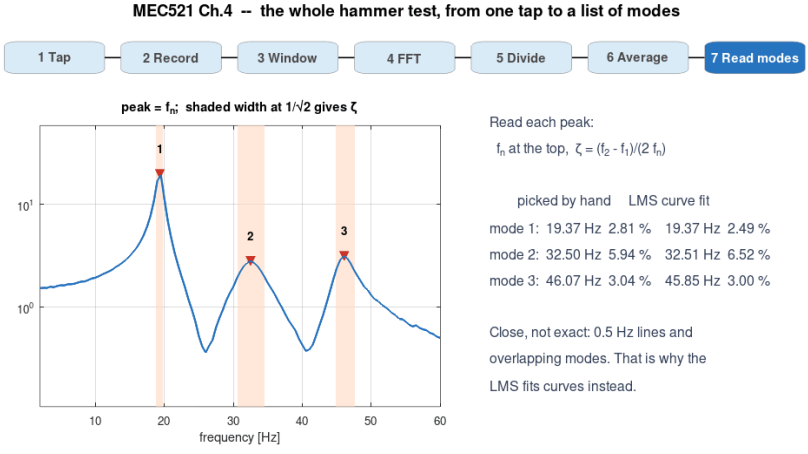
The hammer test, step by step

1. **Support** the structure softly (on cord or foam), or test it as installed.
2. **Tap** with an instrumented hammer: a force sensor in its tip records $F(t)$. An accelerometer records the response $x(t)$.
3. **FFT both**, and divide: $H(f) = X(f)/F(f)$ — the FRF of § 4.1, now measured.
4. **Repeat** 5–10 taps per point and average. Then move to the next point.

Choosing the hammer tip

A **hard** tip gives a short knock that reaches high frequencies; a **soft** tip gives a longer push that only excites low ones. Pick the tip whose force spectrum stays roughly flat — falling no more than about 10 dB — across the band you care about.

The whole test in one animation



One structure with the F-80's first three modes, followed through all seven stages. Pause it at each stage — the next slides take them one at a time. GIF: `fft-anim/MEC521-ch4-pipe.gif`.

What reading the peaks by hand gets you

At the end of the animation the three peaks were read by hand (top of the peak, half-power width) and compared with the LMS curve fit:

	picked by hand		**LMS curve fit**	
mode	f_n [Hz]	ζ [%]	f_n [Hz]	ζ [%]
1	19.37	2.81	19.37	2.49
2	32.50	5.94	32.51	6.52
3	46.07	3.04	45.85	3.00

Frequencies: good. Damping: only roughly. Mode 1's peak is only 0.96 Hz wide — two 0.5 Hz lines — so its width cannot be read finely; mode 3 is pulled by its neighbour.

It is not the noise

Read the *exact* FRF the same way, on the same lines, and mode 1 still reads 2.81 %. The limit is the line spacing and the overlap. That is why commercial software fits curves instead of reading peaks.

Two windows a hammer test needs

Force window
The knock lasts a few milliseconds; after it the force channel holds only noise. Keep the pulse, set the rest to zero.

Exponential window
A lightly damped response may still be ringing when the record ends — leakage. Multiply by e^{-t/τ_w} to force it to die inside the record.

The exponential window adds damping — take it out

Multiplying by e^{-t/τ_w} looks exactly like extra damping:

$$\zeta_{\text{added}} = \frac{1}{2\pi f_n \tau_w} \qquad \text{F4.2.5}$$

A 60 Hz mode with true $\zeta = 0.8\,\%$ and $\tau_w = 0.25\,\text{s}$: $\zeta_{\text{added}} = 1.06\,\%$, so the FRF reads 1.86 %. Software is meant to subtract it — check that it did.

Averaging taps: H_1 and coherence

Each tap is a little different, and the response carries noise. Average over n taps using the products of the spectra:

$$H_1(f) = \frac{\overline{F^* X}}{\overline{F^* F}}, \qquad \gamma^2(f) = \frac{\left|\overline{F^* X}\right|^2}{\overline{F^* F}\;\overline{X^* X}} \quad \in [0,1] \qquad \text{F4.1.4 / F4.1.5}$$

(the bar is the average over the taps). H_1 is the best estimate when the noise is on the response.

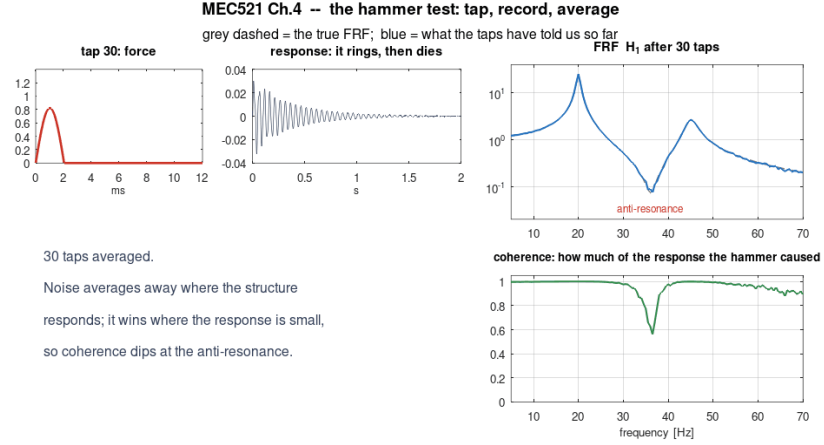
Coherence γ^2 answers one question

At this frequency, how much of the response did the hammer cause? 1 means all of it. It falls where the response is small and noise wins — at **anti-resonances** — and where the hammer put no energy in.

One tap tells you nothing about coherence

With a single tap $\gamma^2 = 1$ at every frequency, whatever the noise. It only means something after several taps.

Watch the FRF build up, tap by tap



A structure with modes at 20 and 45 Hz, tapped 30 times with noise on the response. GIF:

`fft-anim/MEC521-ch4-ham.gif`.

Reading a mode off the FRF

Natural frequency — at the top of the peak (where the phase turns through 90°).

Damping — from the width of the peak at *half power*, where the height has fallen to $1/\sqrt{2}$ of the top:

$$\zeta \approx \frac{f_2 - f_1}{2 f_n} \qquad \text{F2.5.4}$$

How long it rings — the ring-down time constant:

$$\tau = \frac{1}{2\pi\zeta f_n} \qquad \text{F1.4.8} \qquad \text{(after } 5\tau \text{ it is below 1 \%)}$$

Why this is Chapter 2 again

A sharp, narrow peak is a lightly damped mode that rings for a long time. A low, wide peak is a heavily damped one. Two modes closer than about $2\zeta f_n$ blur into one peak.

Worked example — a real drone airframe, mode 1

The F-80 drone, hung on soft cord, was hammer-tested on 28 September 2026 (Siemens LMS, lines every 0.4996 Hz, i.e. a 2.00 s record). Its first mode: $f_n = 19.37\,\text{Hz}$, $\zeta = 2.49\,\%$.

Step 1 — half-power points of that peak

$$f_1 = 18.87\,\text{Hz}, \qquad f_2 = 19.83\,\text{Hz}, \qquad f_2 - f_1 = 0.965\,\text{Hz}$$

Step 2 — damping back from the width

$$\zeta \approx \frac{0.965}{2(19.37)} = \boxed{2.49\,\%}\ \checkmark$$

Step 3 — does it die inside the record?

$$\tau = \frac{1}{2\pi(0.0249)(19.37)} = 0.330\,\text{s}, \qquad 5\tau = 1.65\,\text{s} < 2.00\,\text{s}$$

So no exponential window is needed for this mode — and no added damping to remove.

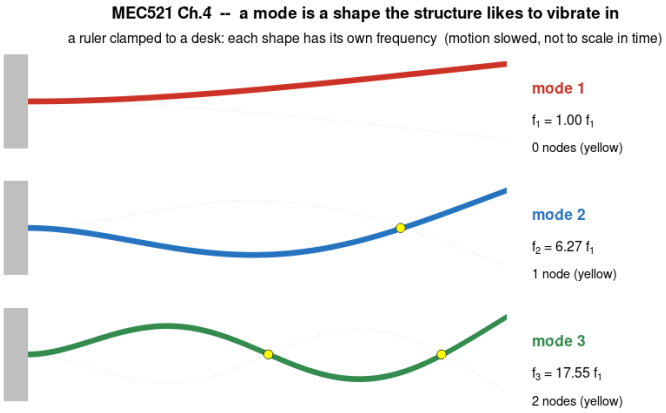
A mode is a shape, not just a frequency

Clamp a ruler to a desk and flick it. It can swing as one curve (mode 1), or with one still point along its length (mode 2), or two (mode 3) — each at its own frequency. The still points are **nodes**.

- to see a shape you must measure **many points**: the size and sign of the response at each one, at f_n, *is* the shape;
- tap or measure **at a node** and that mode is simply missing from your FRF — not because it is not there, but because you were standing on its still point.

For a uniform cantilever, $f_2 = 6.27 f_1$ and $f_3 = 17.55 f_1$; mode 2's node sits at $0.783L$ from the clamp.

Watch three modes at once



Yellow dots are nodes. Motion is slowed and the three are drawn at 1, 2 and 3 swings per loop, not at their true 1 : 6.27 : 17.55. GIF: `fft-anim/MEC521-ch4-modes.gif`.

When peaks overlap, and when there is no hammer

Curve fitting (pole fitting)
Commercial software fits a model with many modes to every measured FRF at once, adding modes until the answers stop moving (a *stabilisation diagram*). The engineer picks the stable ones. That is where a "pole list" comes from.

Operational modal analysis (OMA)
Nobody hammers a bridge. Wind and traffic shake it at all frequencies already; sensors record the response only, and the modes are found from how the sensors move *together*. The Second Penang Bridge carries a monitoring system built on this idea.

The F-80 pole list — 9 modes, curve-fitted

mode	1	2	3	4	5	6	7	8	9
f_n [Hz]	19.37	32.51	45.85	82.29	93.73	104.51	134.66	163.14	249.53
ζ [%]	2.49	6.52	3.00	5.10	2.85	8.22	3.57	7.14	0.67

The one exported FRF (a single point) shows no peak for modes 5–7: that point is near their nodes, or they move sideways.

Hammer test: five ways to get a wrong mode

1. **A double hit.** The hammer bounces; the force spectrum gets deep notches and the FRF is garbage there. Reject the tap.
2. **Too soft a tip.** No force above a few hundred hertz, so no high modes — and no warning, unless you look at the force spectrum.
3. **Exponential-window damping left in.** Every ζ reads high.
4. **Tapping at a node.** The mode is "missing". Test more than one point.
5. **Trusting one tap.** $\gamma^2 = 1$ by definition. Average, then read coherence.

Further reading: P. Avitabile, *Modal Space — Back to Basics* (SEM, free online); A. Brandt, *Noise and Vibration Analysis* (Wiley); D. J. Ewins, *Modal Testing* (RSP). Lab: the *phyphox* app turns a phone into an accelerometer with a live FFT.

Which quantity should you actually measure?

Quantity	Emphasises	Use for
displacement	low frequency	shaft position, clearances
velocity	mid band, flattest	**general severity**
acceleration	high frequency	bearings, gears, impacts

Why velocity is the severity standard

Each differentiation multiplies amplitude by ω. Velocity sits in the middle, so a velocity spectrum is roughly flat across a machine's important band — one number then means the same thing at 10 Hz and 1000 Hz.

Sensors and the mounting trap

Mounting	Usable to	Notes
stud	10 kHz+	the reference; needs a machined face
adhesive	5 kHz	good compromise
magnet	2 kHz	convenient, and the usual culprit
hand-held probe	500 Hz	indicative only

The mount is a spring, and it has its own resonance

A magnet-mounted accelerometer has a mounted resonance around 2 kHz. Above it the readings are the *mount* vibrating, not the machine — and bearing faults live exactly there.

Integration, and why it is not free

$$v_{\mathrm{rms}} = \frac{a_{\mathrm{rms}}}{2\pi f} \qquad \text{F4.3.1}$$

Dividing by f suppresses high frequencies — and **amplifies** low ones. A tiny DC offset or drift in the accelerometer becomes an enormous apparent displacement.

So integrate in the frequency domain, and high-pass first

Integrate the spectrum line by line, after discarding everything below the lowest frequency of interest. Integrating a raw time record twice is how you get a velocity trace that wanders off the screen.

Time-domain descriptors: three numbers before any spectrum

$$\mathrm{CF} = \frac{x_{\mathrm{peak}}}{x_{\mathrm{rms}}} \qquad \text{F4.3.2} \qquad \text{kurtosis} = \frac{\frac{1}{N}\sum (x_i - \bar{x})^4}{\sigma^4} \qquad \text{F4.3.3}$$

Descriptor	Healthy	Rising means
RMS	steady	more energy overall
crest factor	≈ 1.4 for a sinusoid	**impulsiveness**
kurtosis	≈ 3 for Gaussian	spiky, impact-like events

A bearing beginning to spall raises crest factor and kurtosis long before it raises RMS.

Worked example — acceleration to velocity

An accelerometer reads $2.5\,\mathrm{m\,s^{-2}}$ RMS at 50 Hz.

Step 1 — integrate once

$$v_{\mathrm{rms}} = \frac{a_{\mathrm{rms}}}{2\pi f} = \frac{2.5}{2\pi(50)} = 7.96 \times 10^{-3}\,\mathrm{m\,s^{-1}} = 7.96\,\mathrm{mm\,s^{-1}}$$

Step 2 — integrate again

$$d_{\mathrm{rms}} = \frac{v_{\mathrm{rms}}}{2\pi f} = \frac{7.96 \times 10^{-3}}{2\pi(50)} = 25.3\,\mu\mathrm{m}$$

Note the scale

25 µm of movement — invisible, unmeasurable by eye — is $7.96\,\mathrm{mm\,s^{-1}}$ of velocity. Whether that is acceptable is a question for § 4.5, and the answer will surprise you.

Diagnostics: faults have addresses

Every fault excites the machine at a frequency you can *calculate in advance* from the geometry and the shaft speed.

So diagnosis is a matching exercise

1. Compute the expected frequencies — shaft orders, gear mesh, bearing defects.
2. Measure the spectrum.
3. See which predicted lines are present, and how large.

$$\mathrm{GMF} = N_{\text{teeth}} \times f_{\text{shaft}} \qquad \text{F4.4.1}$$

Which is why the shaft speed is the single most important number on the sheet. Without it, nothing can be attributed.

Shaft-order faults

Faults tied to the shaft appear at integer multiples of f_r — the **orders**.

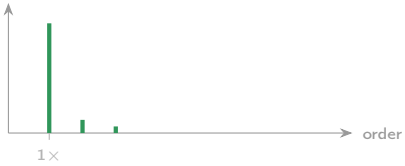
Fault	Signature	Distinguishing feature
imbalance	$1\times$ dominant	radial, steady phase
misalignment	$2\times$ strong	often high axial
looseness	many harmonics	$\frac{1}{2}\times$ subharmonics
bent shaft	$1\times$	axial, phase differs across the coupling

$1\times$ alone does not mean imbalance

A bent shaft gives $1\times$ too. The direction of the vibration and the phase relationship are what separate them — amplitude alone never does.

Four signatures worth recognising on sight

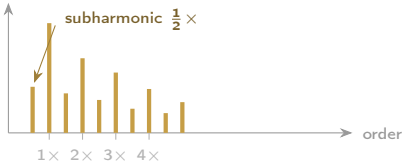
Unbalance
dominant 1×, little else



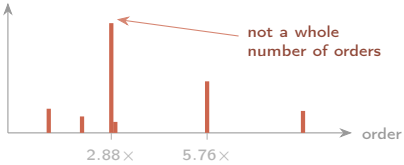
Misalignment
2× comparable with or above 1×



Looseness
many harmonics, often $\frac{1}{2}\times$ too



Bearing defect
non-integer order, with harmonics



Bearing defect frequencies

$$\mathrm{FTF} = \frac{f_r}{2}\left(1 - \frac{d}{D}\cos\phi\right) \qquad \text{F4.4.2}$$

$$\mathrm{BPFO} = \frac{nf_r}{2}\left(1 - \frac{d}{D}\cos\phi\right) \qquad \mathrm{BPFI} = \frac{nf_r}{2}\left(1 + \frac{d}{D}\cos\phi\right)$$

F4.4.3 F4.4.4

Two identities that check your arithmetic for free

$$n\,\mathrm{FTF} = \mathrm{BPFO}, \qquad \mathrm{BPFO} + \mathrm{BPFI} = nf_r \qquad \text{F4.4.6}$$

Why bearing frequencies are never whole numbers

$$\frac{\text{defect frequency}}{f_r} \notin \mathbb{Z} \qquad \text{F4.4.7}$$

The rolling elements *slip* slightly; the cage rate is not a rational fraction of shaft speed.

So a bearing line lands **between** the shaft orders.

This is a feature, not a nuisance

It is exactly what lets you tell a bearing fault from a shaft fault. A line at $3.06\times$ is a bearing; a line at exactly $3\times$ is not.

But only if your resolution can see the difference

BPFO at 89.25 Hz and $3f_r$ at 87.5 Hz are 1.75 Hz apart. At $\Delta f = 2$ Hz they are one line — and that line looks like misalignment.

Worked example — bearing frequencies (1 of 2)

A bearing with $n = 9$ rolling elements, $d/D = 0.32$, contact angle $\phi = 0$, on a shaft at 1750 rev/min.

Step 1 — shaft rate first, always

$$f_r = \frac{1750}{60} = 29.17\,\text{Hz}$$

Step 2 — the four defect frequencies

$$\text{FTF} = \tfrac{1}{2}(0.68)(29.17) = 9.92\,\text{Hz}, \qquad \text{BPFO} = \tfrac{9}{2}(0.68)(29.17) = 89.25\,\text{Hz}$$

$$\text{BPFI} = \tfrac{9}{2}(1.32)(29.17) = 173.25\,\text{Hz}, \qquad \text{BSF} = \frac{1 - 0.32^2}{2(0.32)}(29.17) = 40.91\,\text{Hz}$$

Step 3 — check with the two identities

$$9 \times 9.92 = 89.25 \;\checkmark \qquad 89.25 + 173.25 = 262.5 = 9 \times 29.17 \;\checkmark$$

Worked example — bearing frequencies (2 of 2)

Step 4 — will the measurement even resolve it?

The third shaft harmonic sits at $3f_r = 87.5\,\text{Hz}$, only $1.75\,\text{Hz}$ below BPFO.

To place at least two spectral lines in that gap:

$$\Delta f \leq 0.875\,\text{Hz} \qquad \Longrightarrow \qquad N \geq \frac{f_s}{0.875} = \frac{5120}{0.875} = 5851$$

$N = 8192$	$\Delta f = 0.625\,\text{Hz}$	sufficient ✓
$N = 4096$	$\Delta f = 1.25\,\text{Hz}$	the two merge — **no**

Getting this wrong invents a fault

At $N = 4096$ the bearing line and the shaft harmonic become one peak at $\approx 3\times$ — and the report says "misalignment".

Envelope analysis

A spalled bearing produces tiny, sharp *impacts* — repeated at BPFO but each one containing almost no energy at BPFO itself.

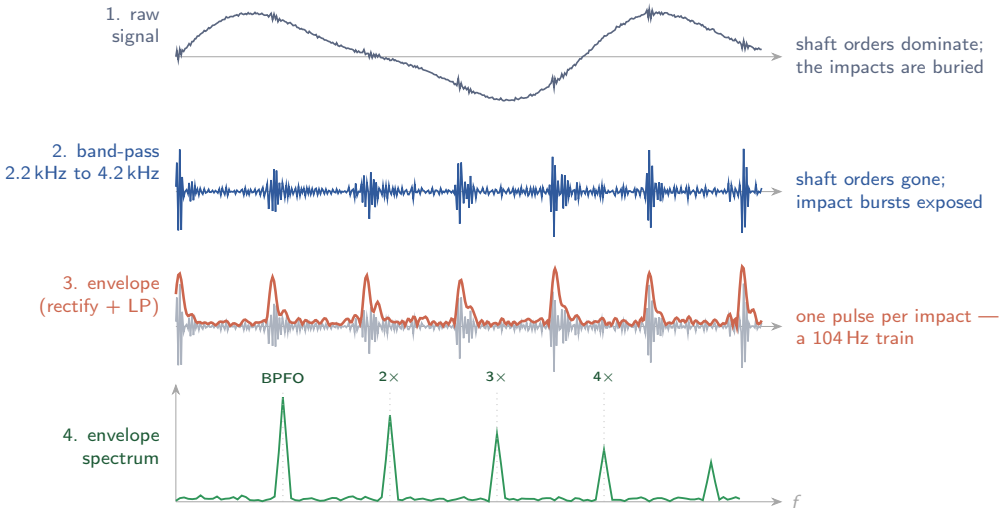
The trick

The impacts ring the structure at a high resonance. That high-frequency carrier is **amplitude-modulated** at the defect rate — so demodulate it:

1. band-pass around the structural resonance;
2. rectify and smooth, to extract the envelope;
3. take the spectrum *of the envelope*.

BPFO appears cleanly in the envelope spectrum, where it was invisible in the ordinary one.

Envelope analysis, step by step



When to reach for the envelope

Ordinary spectrum	Envelope spectrum
imbalance, misalignment, looseness	early bearing faults
gear mesh	gear tooth cracks
strong, low-frequency energy	weak, repetitive impacts

The band-pass choice is the whole method

Pick the band around the *structural* resonance that the impacts excite — typically 2 kHz to 10 kHz. Choose it wrongly and the envelope spectrum is empty, and the bearing looks healthy.

Severity: ISO 10816

Machine class sets the thresholds. Velocity RMS is the measured quantity.

Zone	Class I	Class II	Meaning
A	< 0.71	< 1.12	newly commissioned
B	< 1.8	< 2.80	acceptable for long-term running
C	< 4.5	< 7.10	unsatisfactory; plan an intervention
D	> 4.5	> 7.10	**capable of causing damage**

($\mathrm{mm\,s^{-1}}$ RMS) F4.5.1 F4.5.2

Class I is under 15 kW; Class II is medium machines on rigid foundations.

Worked example — severity assessment

The reading from § 4.3 — 2.5 m s^{-2} RMS at 50 Hz — was taken on a 7.5 kW pump.

Step 1 — classify the machine 7.5 kW is below 15 kW, so **Class I**.

Step 2 — convert to the standard's quantity

$$v_{\mathrm{rms}} = \frac{2.5}{2\pi(50)} = 7.96\,\mathrm{mm\,s^{-1}}$$

Step 3 — read the zone

$$7.96 > 4.5 \qquad \Longrightarrow \qquad \boxed{\text{Zone D}}$$

What Zone D actually means

Not "schedule it for the next shutdown". Zone D on a Class I machine means the vibration is **capable of causing damage** — investigate now and plan an intervention.

Monitoring is not protection

Condition monitoring

ISO 10816. Periodic, diagnostic, informs *maintenance planning*. Answers: what is wrong, and how urgent?

Machinery protection

API 670. Continuous, automatic, *trips the machine*. Answers one question: shut down, yes or no?

Different systems, different sensors, different answers

A protection system is deliberately crude — it must never trip spuriously and must never fail to trip. It diagnoses nothing.

$$t = \frac{\ln\left(v_{\text{limit}} / v_{\text{now}}\right)}{r} \qquad \text{F4.5.3}$$

— how long until the trend reaches the limit

From measurement to recommendation

A report that stops at "the vibration is $7.96\,\mathrm{mm\,s^{-1}}$" has done half the job.

1. **What** — the measured level and the zone.
2. **Where** — which frequency, and which fault it matches.
3. **How confident** — does a second identity or a second sensor agree?
4. **How urgent** — from the trend, not from the single reading.
5. **What to do** — and what it costs to do nothing.

This is the structure your Mini Project report is marked against.

Machine learning for anomaly detection

When there is no physical model, learn what *normal* looks like and flag departures from it.

Approach	Needs	Gives
threshold on a feature	one healthy baseline	an alarm
clustering / novelty	healthy data only	"this is unusual"
supervised classifier	labelled *faults*	a named fault

Labelled fault data is the thing nobody has

Machines are not run to destruction to build a dataset. So most practical condition monitoring is novelty detection — and it reports *that* something changed, never *what*.

Case study: prediction against measurement

A tilt-rotor VTOL UAV. FE modal analysis extracted ten modes, the highest at 93.95 Hz. It then flew **five instrumented hovers over three months**.

Dominant measured peaks, motors loaded	168 Hz to 221 Hz
Front two-blade pass, from the tracked fundamental	219 Hz
Highest extracted FE mode	93.95 Hz

The model is not wrong. It is silent about the band the aircraft actually vibrates in. Ten modes is a solver default, not an engineering decision — and the excitation envelope was built on *shaft* rate. A two-blade propeller forces at *twice* shaft rate, so the real forcing line was never inside the envelope, nor the extraction, nor the harmonic sweep sized to match them. **One factor of blade count, three ceilings.**

The peak could not be attributed

Attribution needs a variable the *excitation* depends on and the *structure* does not. Rotor speed is that variable — and it was never logged. Commanded throttle is the substitute every autopilot records. Across the five flights its correlation with the tracked blade-pass frequency:

flight	1	2	3	4	5
r	-0.738	$+0.101$	-0.325	-0.041	-0.156

Three of the correlations in this campaign are individually significant — and they disagree in sign. That is not a weak relationship. It is the absence of one — and a proxy that fails this way still returns a confident number on any single dataset.

A modal test with none of the equipment



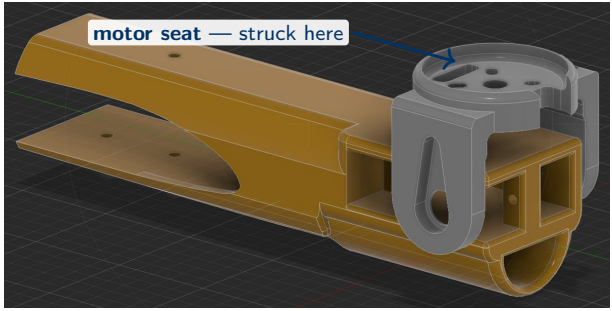
No shaker. No force transducer. No accelerometer of its own.

- Elastic cord on a keyboard stand
- A mallet, force unmeasured
- The aircraft's *own* autopilot IMU

Thirty minutes. It identified what three months of flight data could not.

What you give up is fixed by what you left out: no input force means no FRF, no mode shapes, and no damping. Frequencies only — and the report must say so.

Strike the load path, not the convenient spot



struck at	response	
Motor seat	$\times 71$	the nacelle modes rang
Nose battery tray	$\times 10$	nothing usable

If it does not ring, suspect the location before the structure.

The rig answers loudest, and it is not the structure

A suspension is a soft spring, and springs have a frequency:

$$f_s = \frac{1}{2\pi}\sqrt{\frac{g}{\delta}}, \qquad \delta = \text{static sag}$$

Measured sag on that stand: just over 5 mm, predicting 7.0 Hz. Measured response: a mode at 7.96 Hz, answering $\times 10\,513$ above the floor — louder than anything structural in the test. **Measure the sag before the first strike.** It turns a mystery peak into an arithmetic check.

Rule of thumb: f_s five times below the lowest mode of interest. 4 mm of sag buys 8 Hz; 27 mm buys 3 Hz.

The result, and the control that makes it evidence

Impacts on a motor mount rang six modes between 160 and 193 Hz. The flight peaks in that band sat a mean of 1.49 Hz away, against 6.31 Hz by chance.

	mean $\lvert\Delta\rvert$	p
Family under test	1.49 Hz	0.0001
Control — a known rotor order	6.33 Hz	0.53

Six modes across 100 Hz put a ± 3 Hz window over half the band. About half of *any* frequency list would appear to match. **So the same test is run on a family that must NOT match.**

It does not. A test that has never been shown to fail is not evidence when it passes.

Three ways to lose marks here

1. **Quoting a spectrum without its resolution.**
   Δf decides whether two lines are separable. A fault identification is worthless without it.

2. **Comparing acceleration with a velocity limit.**
   ISO 10816 is in $\mathrm{mm\,s^{-1}}$ RMS. Integrate first, then compare.

3. **Naming a fault from amplitude alone.**
   $1\times$ is imbalance *or* a bent shaft. Say what additional evidence would separate them.

And the one that invalidates everything

No shaft speed, no diagnosis. Every defect frequency is proportional to f_r.

What the assessment will ask

Mini Project	Wk 14, 10 %	CLO3 (C5) — you are given a real dataset and must *evaluate* it.
Final	60 %	CLO3 parts ask you to judge, recommend and justify.

What a full-mark answer looks like

- the acquisition settings stated and justified;
- predicted frequencies computed *before* the spectrum is read;
- at least one independent check — an identity, or a second sensor;
- a severity judgement against a named standard and class;
- a recommendation, with its urgency and its cost of inaction.

What to take away

1. $\Delta f = 1/T$. Resolution comes from record *length*, not sample rate.
2. Aliasing is the only error you cannot undo — filter before sampling.
3. Bearing frequencies land *between* shaft orders. That is what identifies them.
4. Severity is judged in velocity RMS, against a machine class, using a named standard.

Industry alignment

Energy, Aerospace, Railway Engineering — vibration data-driven maintenance to detect misalignment, imbalance in bearings and shafts, wheel/track defects. Handbook § 4.1–4.5.

Technology integration

AI for predictive maintenance; machine learning for anomaly detection.

Lab integration

Use a basic ML model to compare normal and faulty signals for anomaly detection; time-domain vibration data to obtain the FRF.

The three you will reach for most: F4.4.3, F4.2.1, F4.5.2.
Everything else is on *the short list* at the front of Appendix A — one page, ordered by what you are looking at.

Tutorial: *Condition Monitoring and Diagnostics* (30 kW process pump). The rehearsal for the Mini Project.