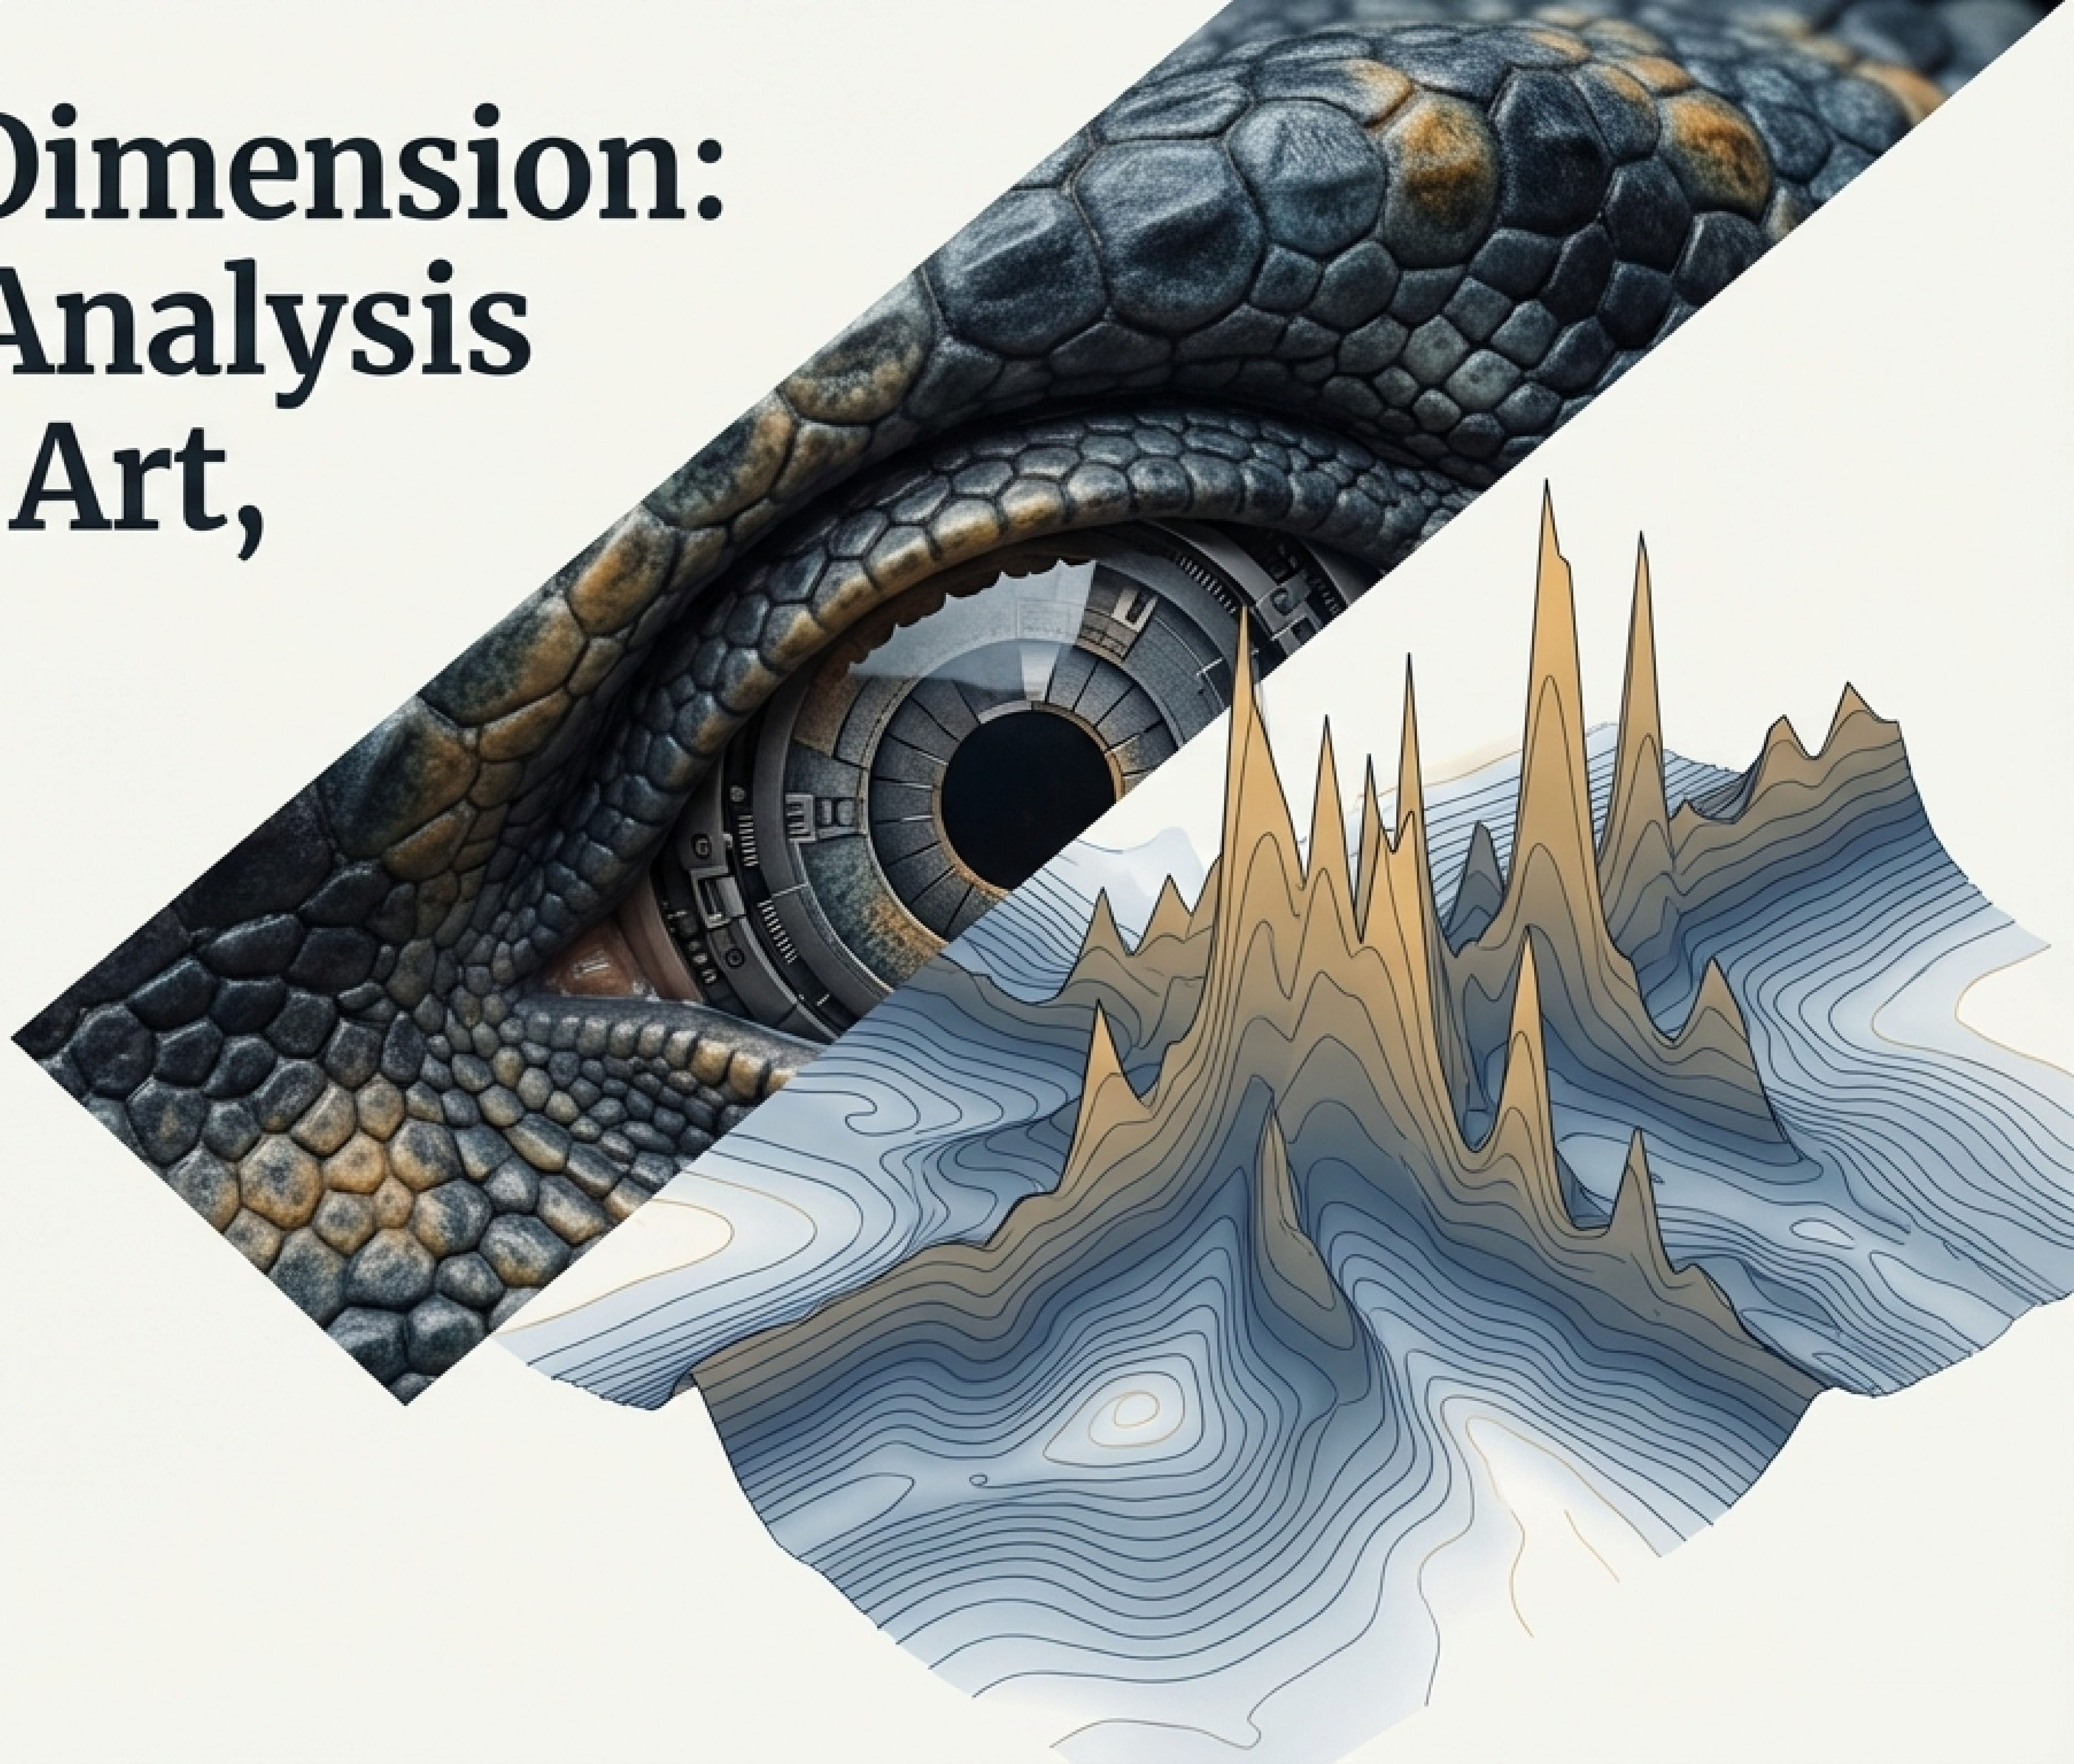


The Invisible Dimension: How Spectral Analysis Shapes Audio, Art, and AI

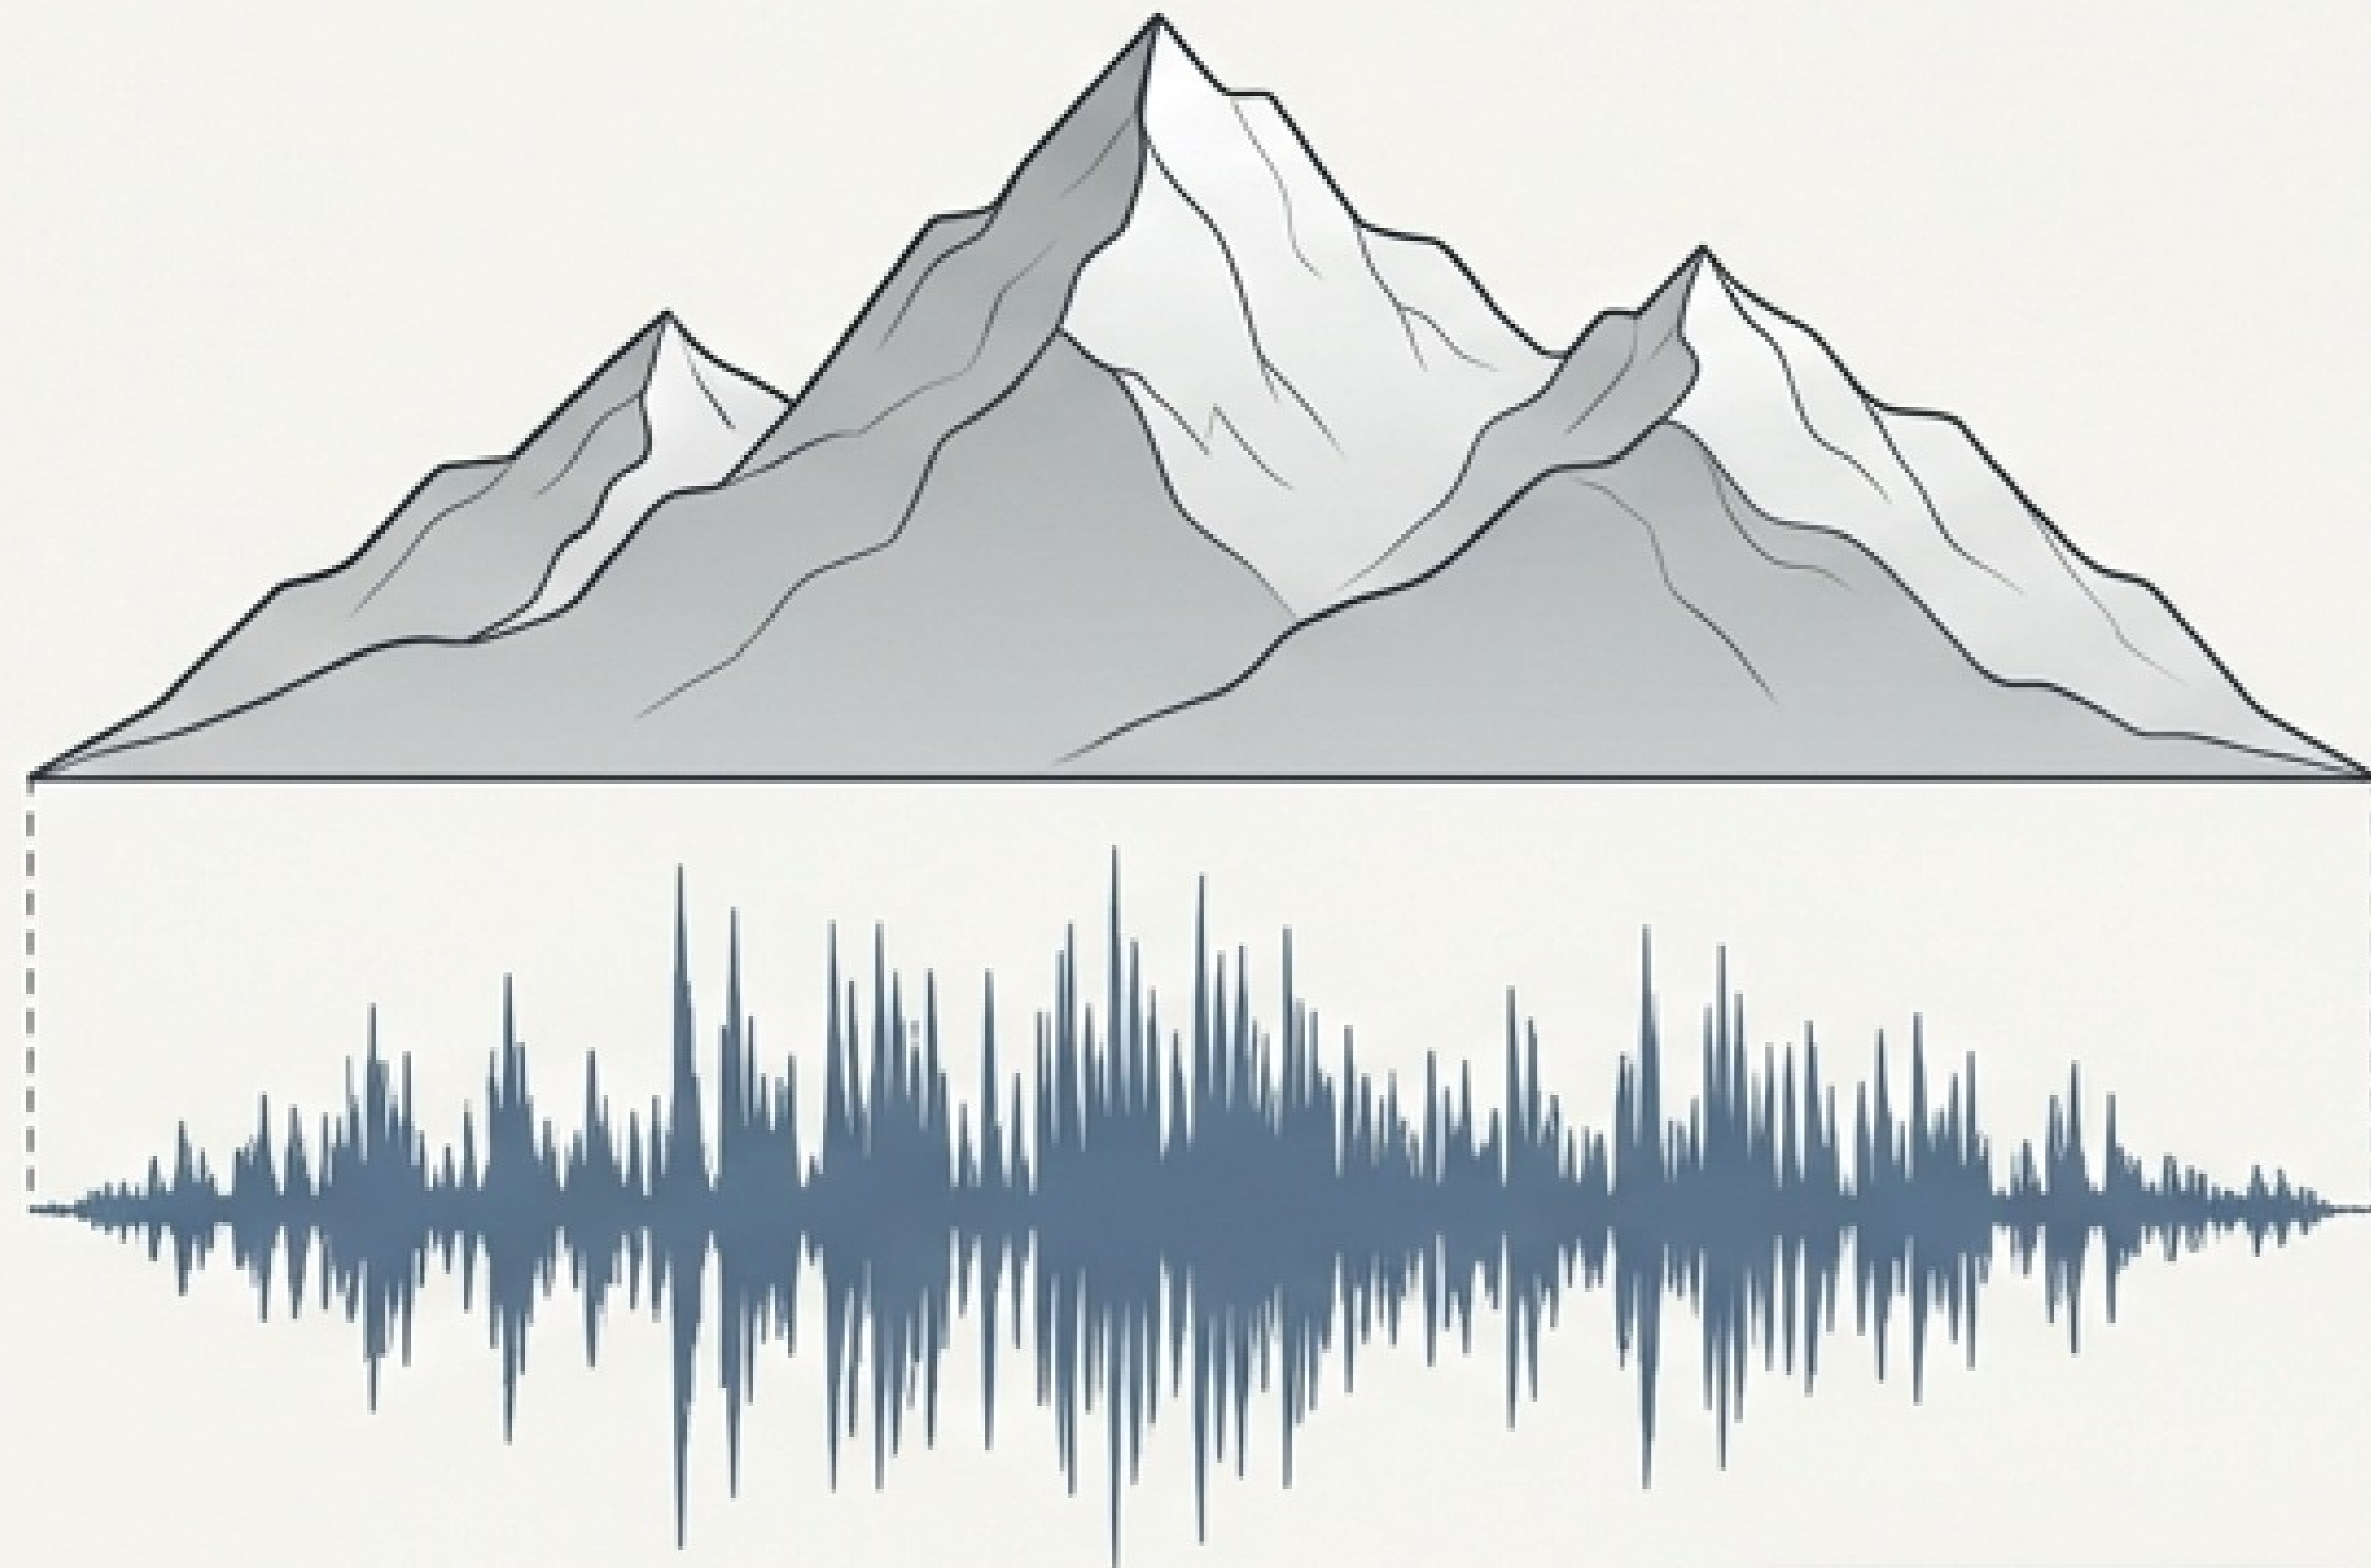


A Visual Synthesis of Frequency
Domain Applications.

Translating Reality into Waves

The Fourier Principle. Any well-behaved function can be represented by a **superposition** of sinusoidal waves.

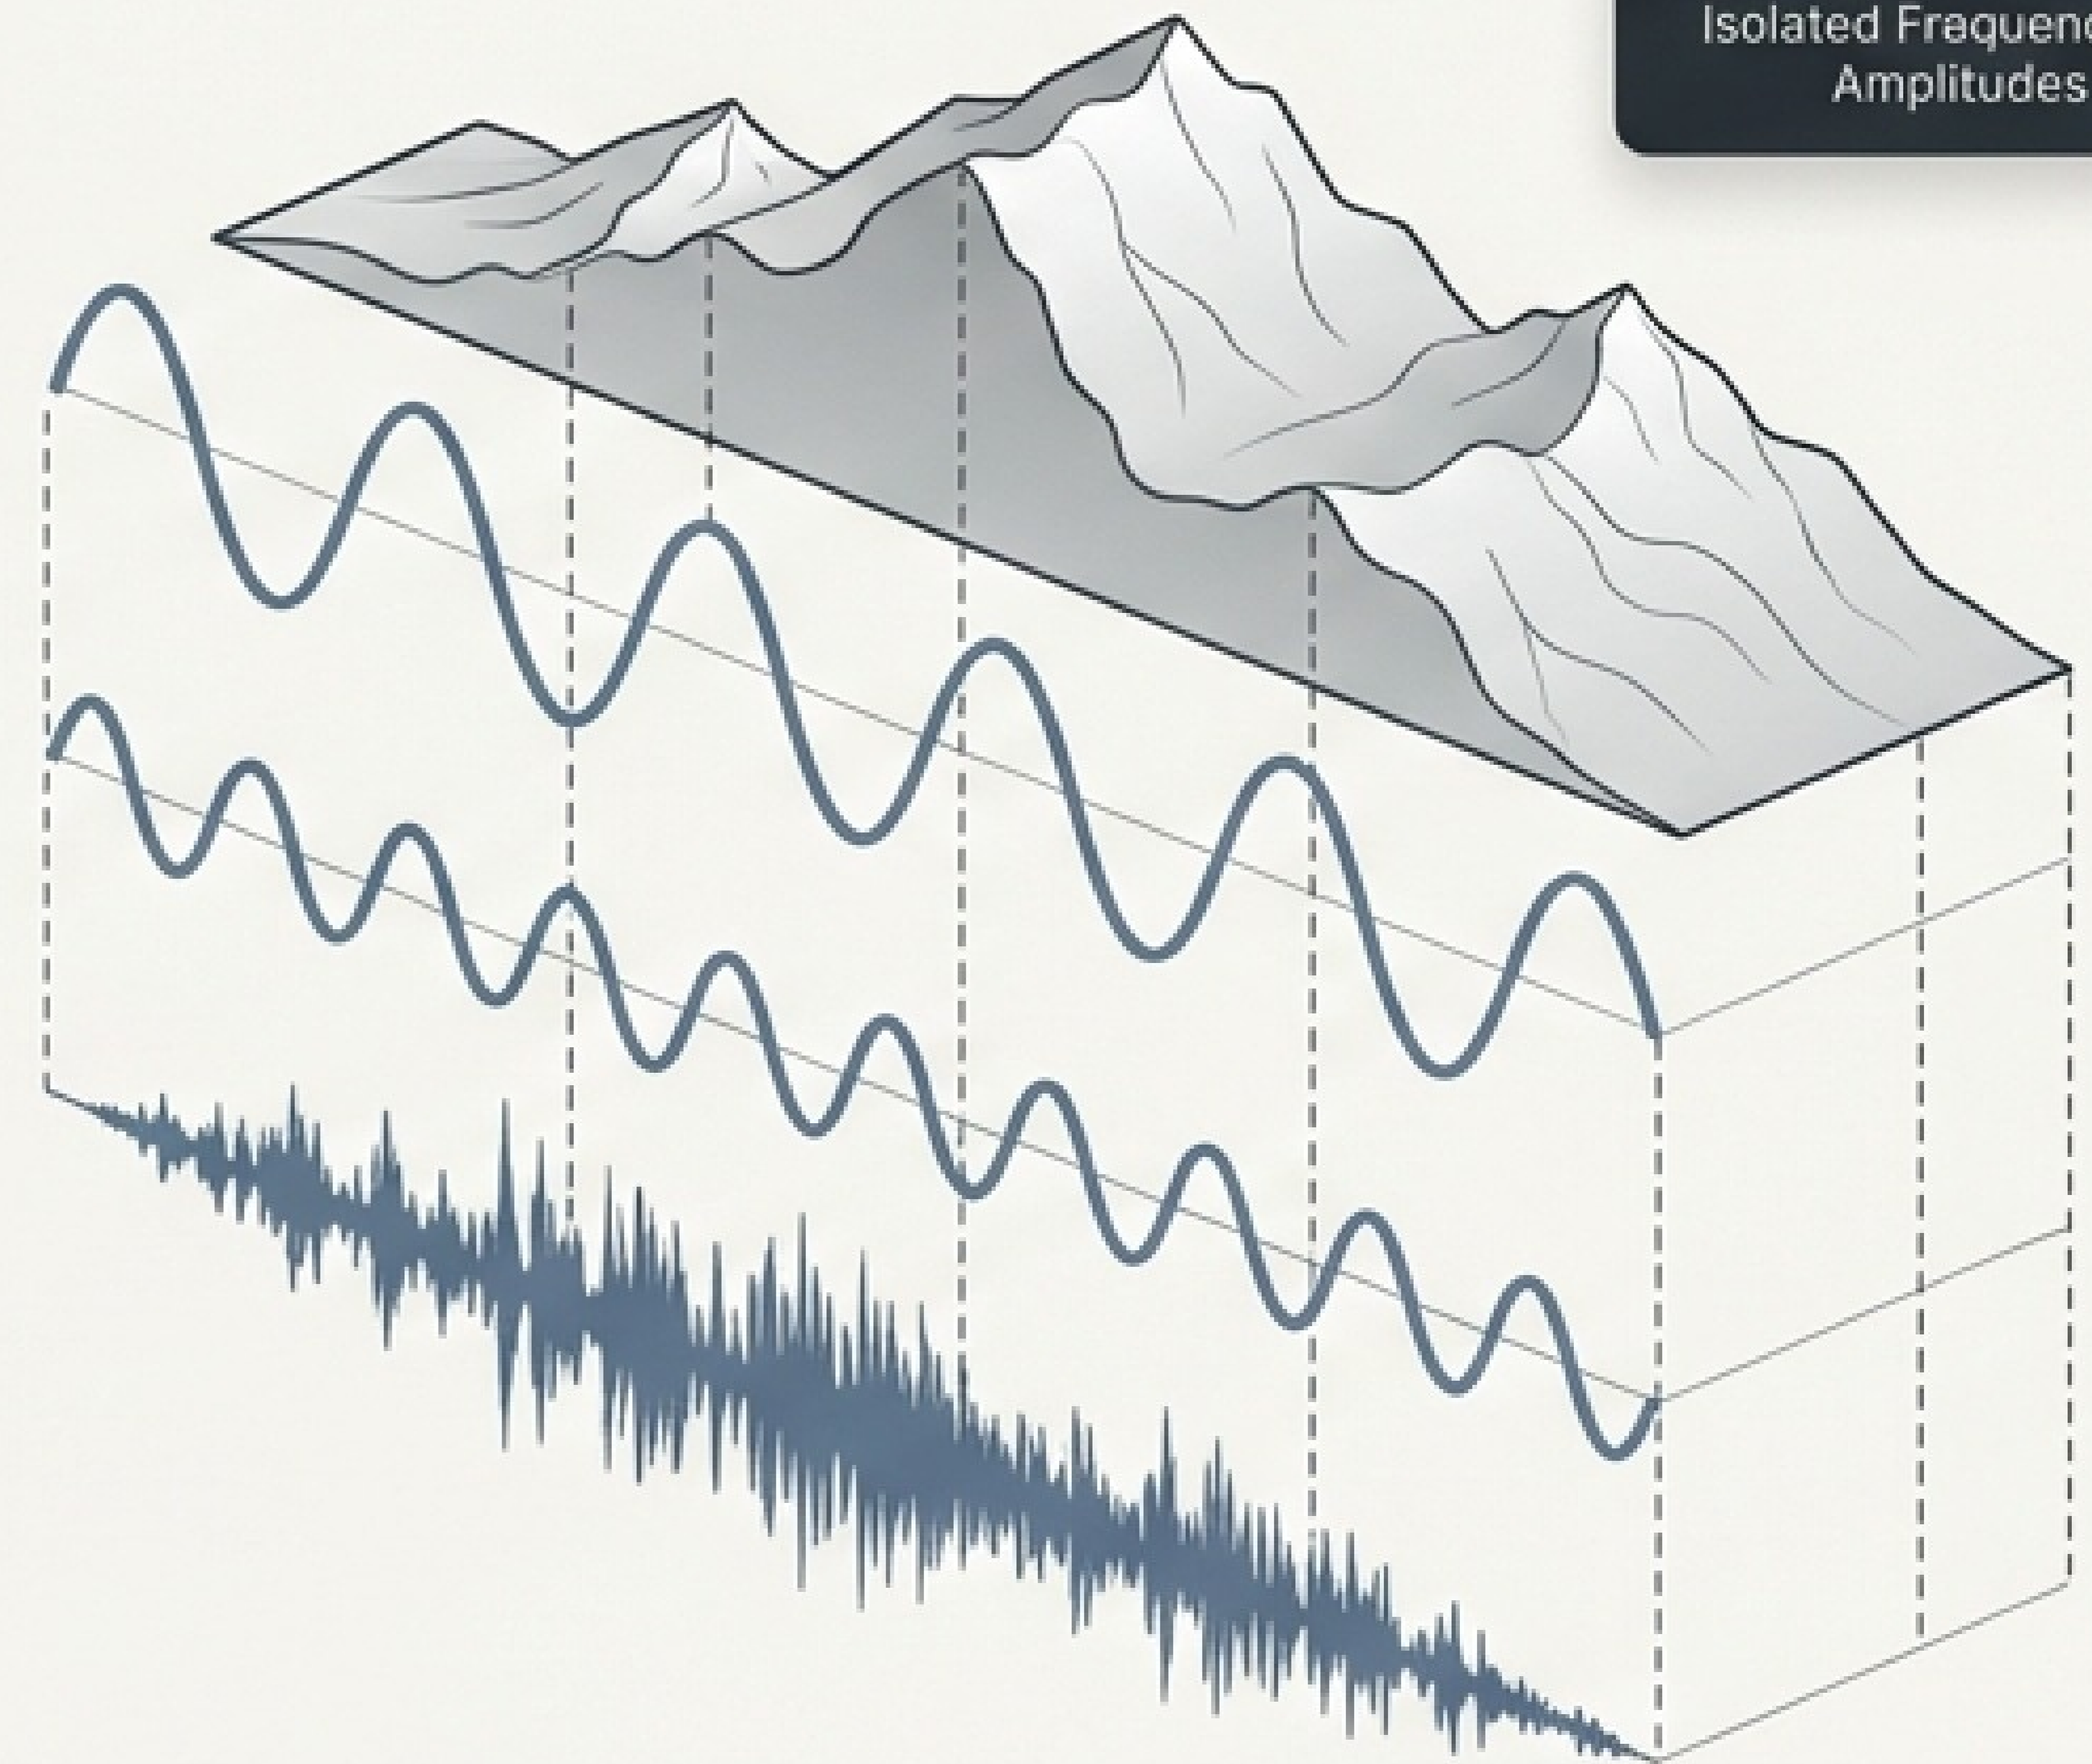
Spatial/Time Domain



Complex Signal
Superposition of Waves




Frequency Domain



Component Waves
Isolated Frequencies &
Amplitudes

The Shift. Moving from spatial coordinates (X, Y) or time to frequencies and amplitudes allows us to isolate, manipulate, and rebuild reality.

The Engines of Analysis

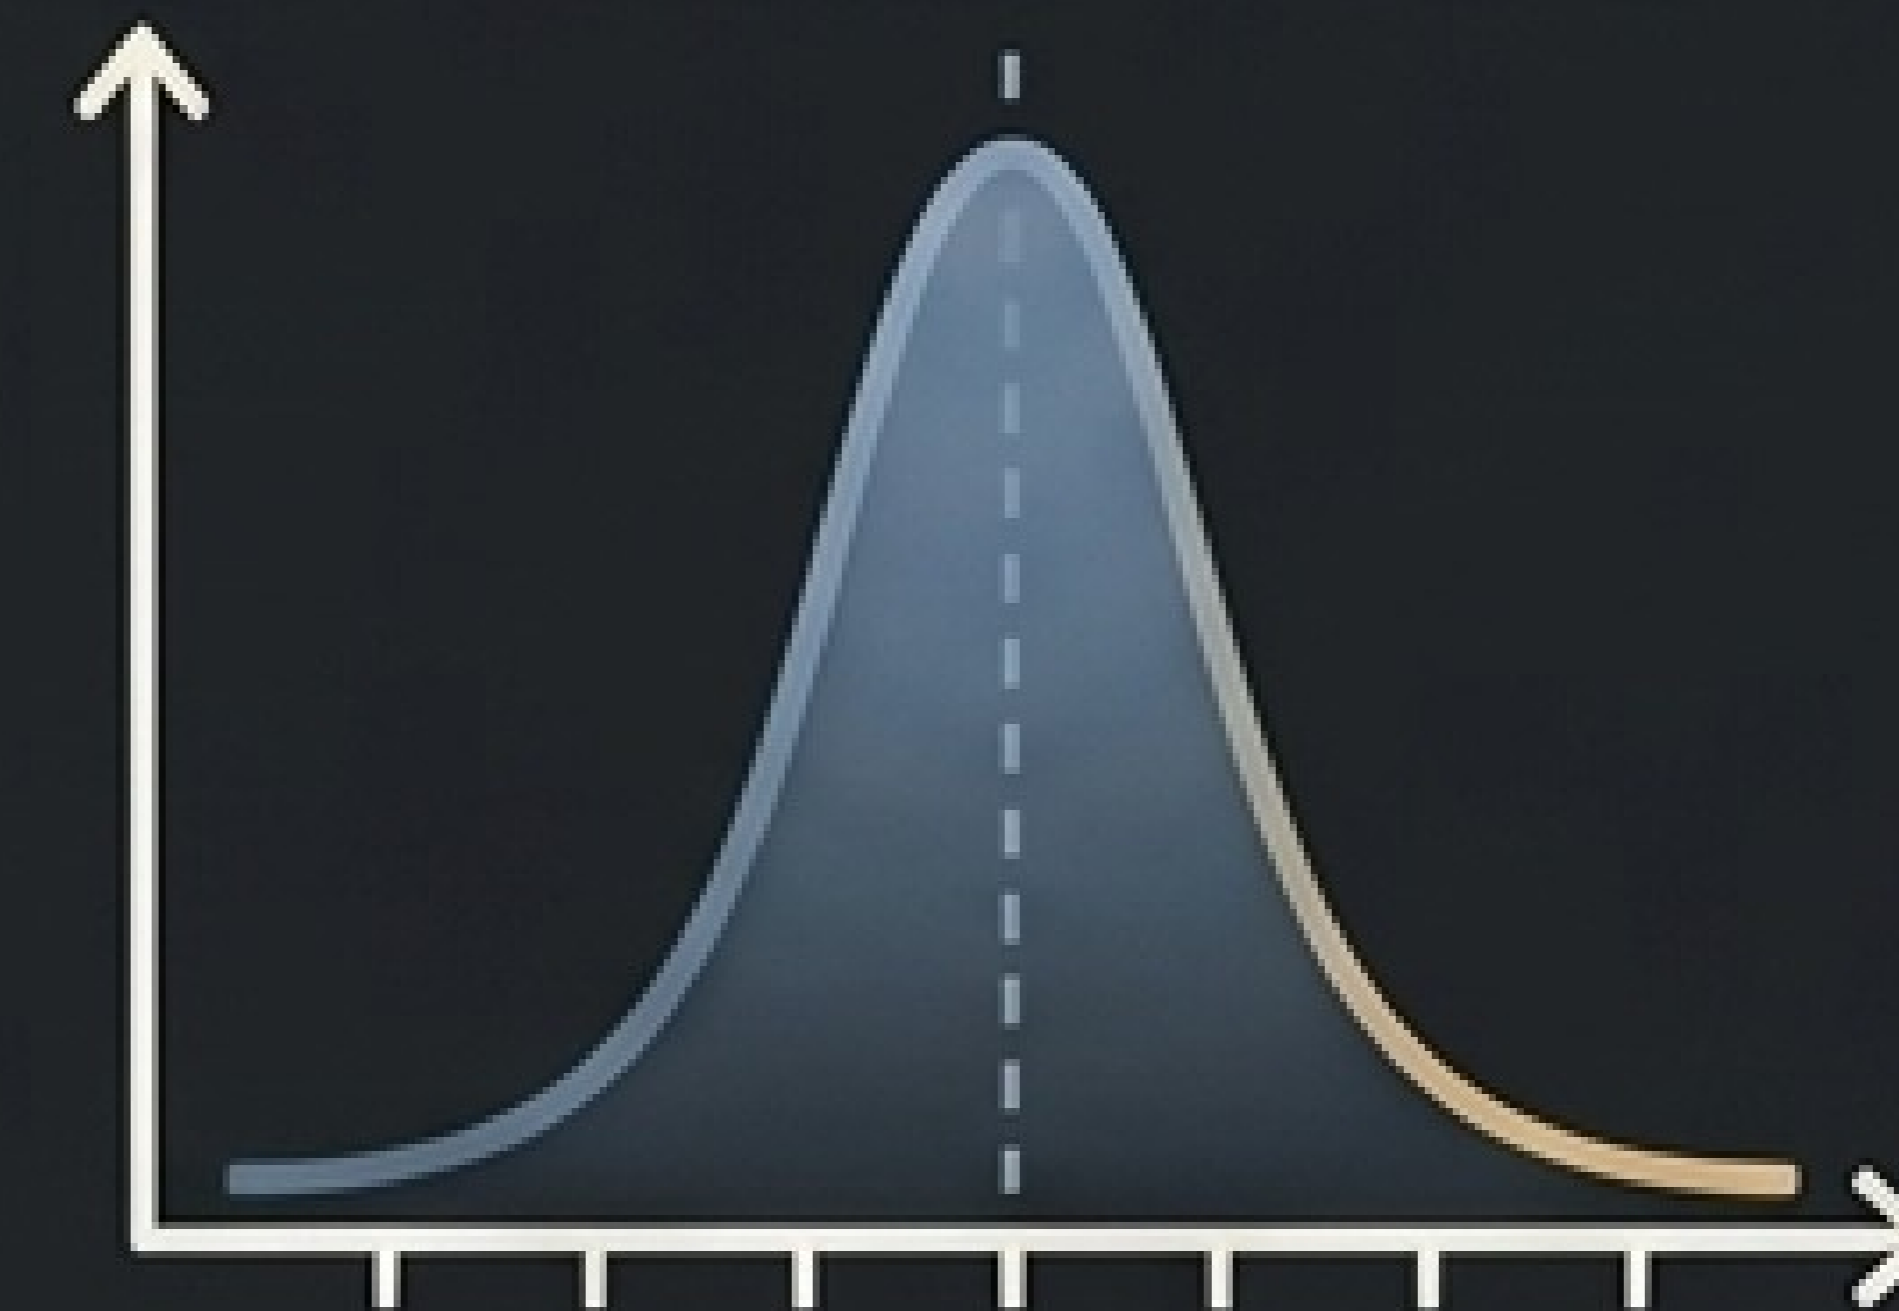


Fast Fourier Transform (FFT)

Function: Converts time-domain data into the complex frequency domain.

Efficiency: Reduces discrete computation complexity from $O(N^2)$ to $O(N \log N)$.

Output: Amplitude, phase, and frequency components.




Power Spectral Density (PSD)

Function: Quantifies how a signal's power is distributed across frequencies.

Efficiency: Normalised to a 1 Hz bandwidth for consistent power values across varying bandwidths.

Output: A real-valued statistical representation (power per unit frequency).

Decoding the 2D Fourier Transform




Original Image



Magnitude Image

Magnitude

Stores colour and power. Large features concentrate near the centre; small details scatter outward.




Phase Image


Phase

Stores position and structure. Ranges from $-\pi$ to $+\pi$, dictating exactly where edges and shapes align.

Pillar 1: Shaping Sound



Engineering a Prehistoric Roar



1. **Source:** Asian elephant rumble (22 Hz fundamental).

2. **Time & Pitch:** Down-shift -1.8 octaves (-214 cents) and stretch duration by 125%.

3. **Spectral Sculpting:** Apply 48 dB/octave low-pass filtering at 150 Hz to isolate subterranean rumble.

4. **Spatial Authenticity:** Add convolution reverb with a 2.4-second decay tail (9,000 m² cave impulse response).

Technical Specifications

Target SPL: 89-94 dB @ 15m

Sync Latency: <2.8 ms (oscilloscope verified)

Timing Tolerance: ±0.5 ms with 120 fps strobe sync

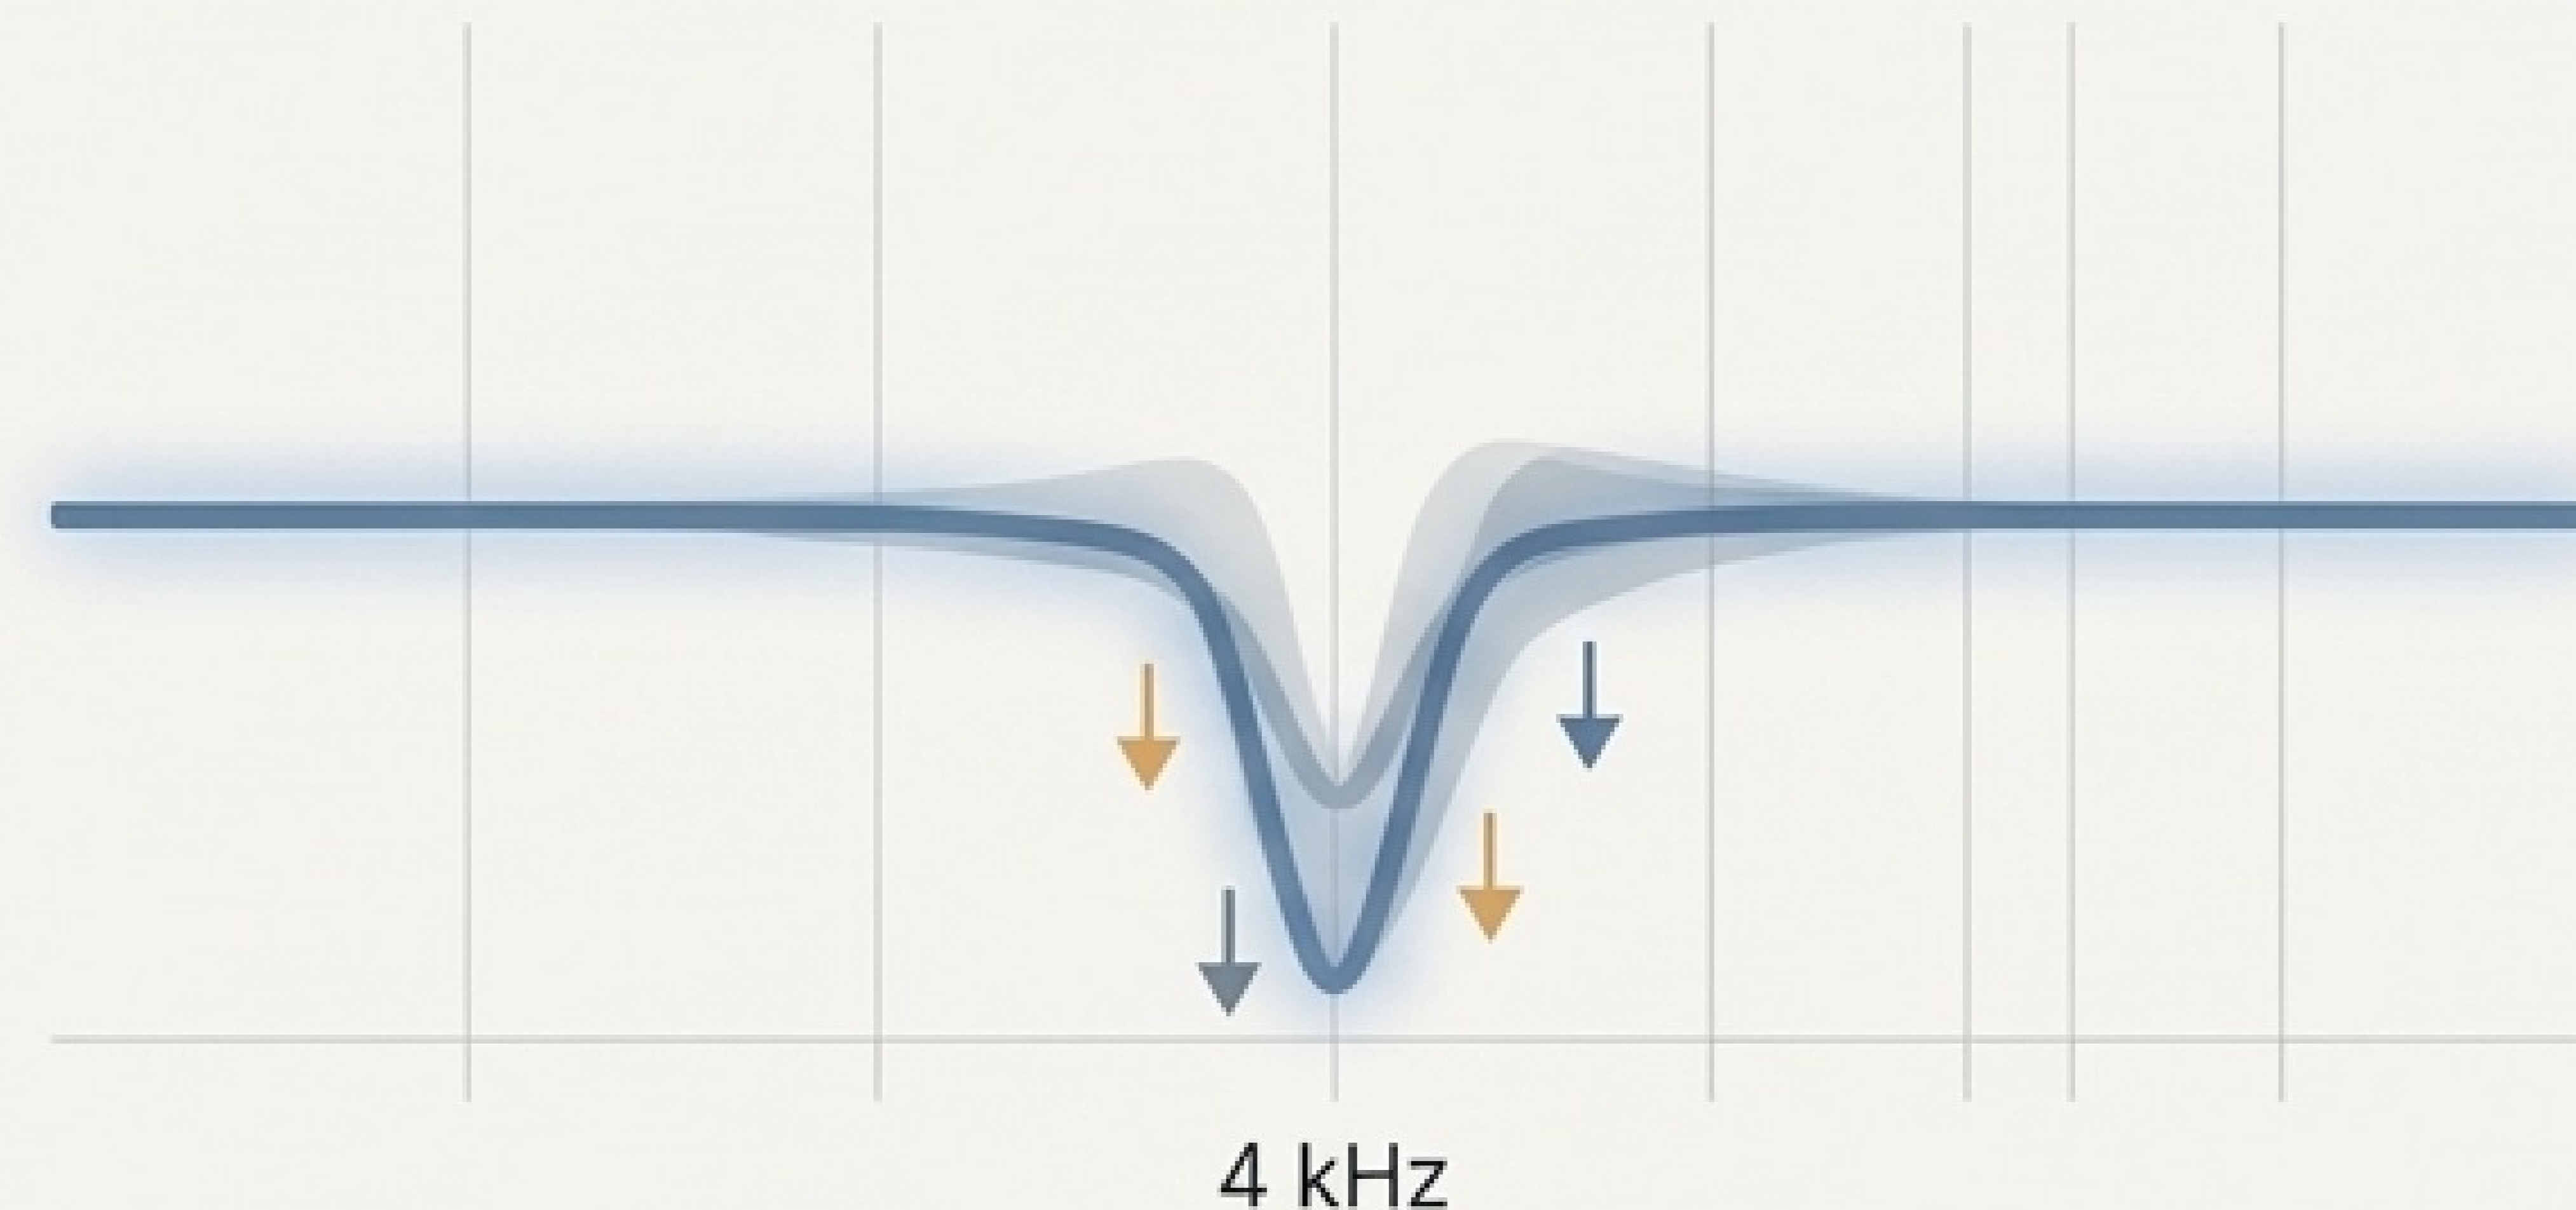
Advanced Tonal Control: Dynamic EQ



STATIC EQ

The Problem with Static

Cutting 3-5 kHz to tame harsh vocal sibilance permanently dulls the entire mix.



DYNAMIC EQ

The Dynamic Solution

Plugins like SSL X-DynEQ (24 independent nodes) or FabFilter (Spectral Dynamics) analyse the material and adapt response times automatically.


Result

Surgical control of transient spikes (e.g., cymbal overtones) while leaving musical harmonic content completely untouched.

Pillar 2: Sculpting Light



Photoshop's Frequency Separation



Layer 1: Low Frequency (Shadows & Colour)

Achieved via Gaussian Blur (e.g., 5px radius). Stores tone and lighting.

Layer 2: High Frequency (Texture & Detail)

Achieved via Apply Image (Target: Low Freq layer, Blending: Subtract, Scale: 2, Offset: 128). Set blend mode to Linear Light.


By separating spatial details into frequency bands, retouchers can smooth skin tones without destroying pores, or remove blemishes without smudging underlying shadows.

Spatial Filtering via Fourier Transforms

Original




The Spectrum Masks




Low-Pass Filtering

A circular mask over the spectrum centre. Removes high frequencies to create perfect convolution blurring.



High-Pass Filtering

The inverse mask. Removes the centre DC value to extract sharp edge detection.




Notch Filtering

Manually masking specific 'star-like' dots in the spectrum to instantly delete repeating spatial noise (like dither or grid patterns) without blurring the entire image.

Resulting Cleaned Image



Pillar 3: Next- Generation AI & Sensors



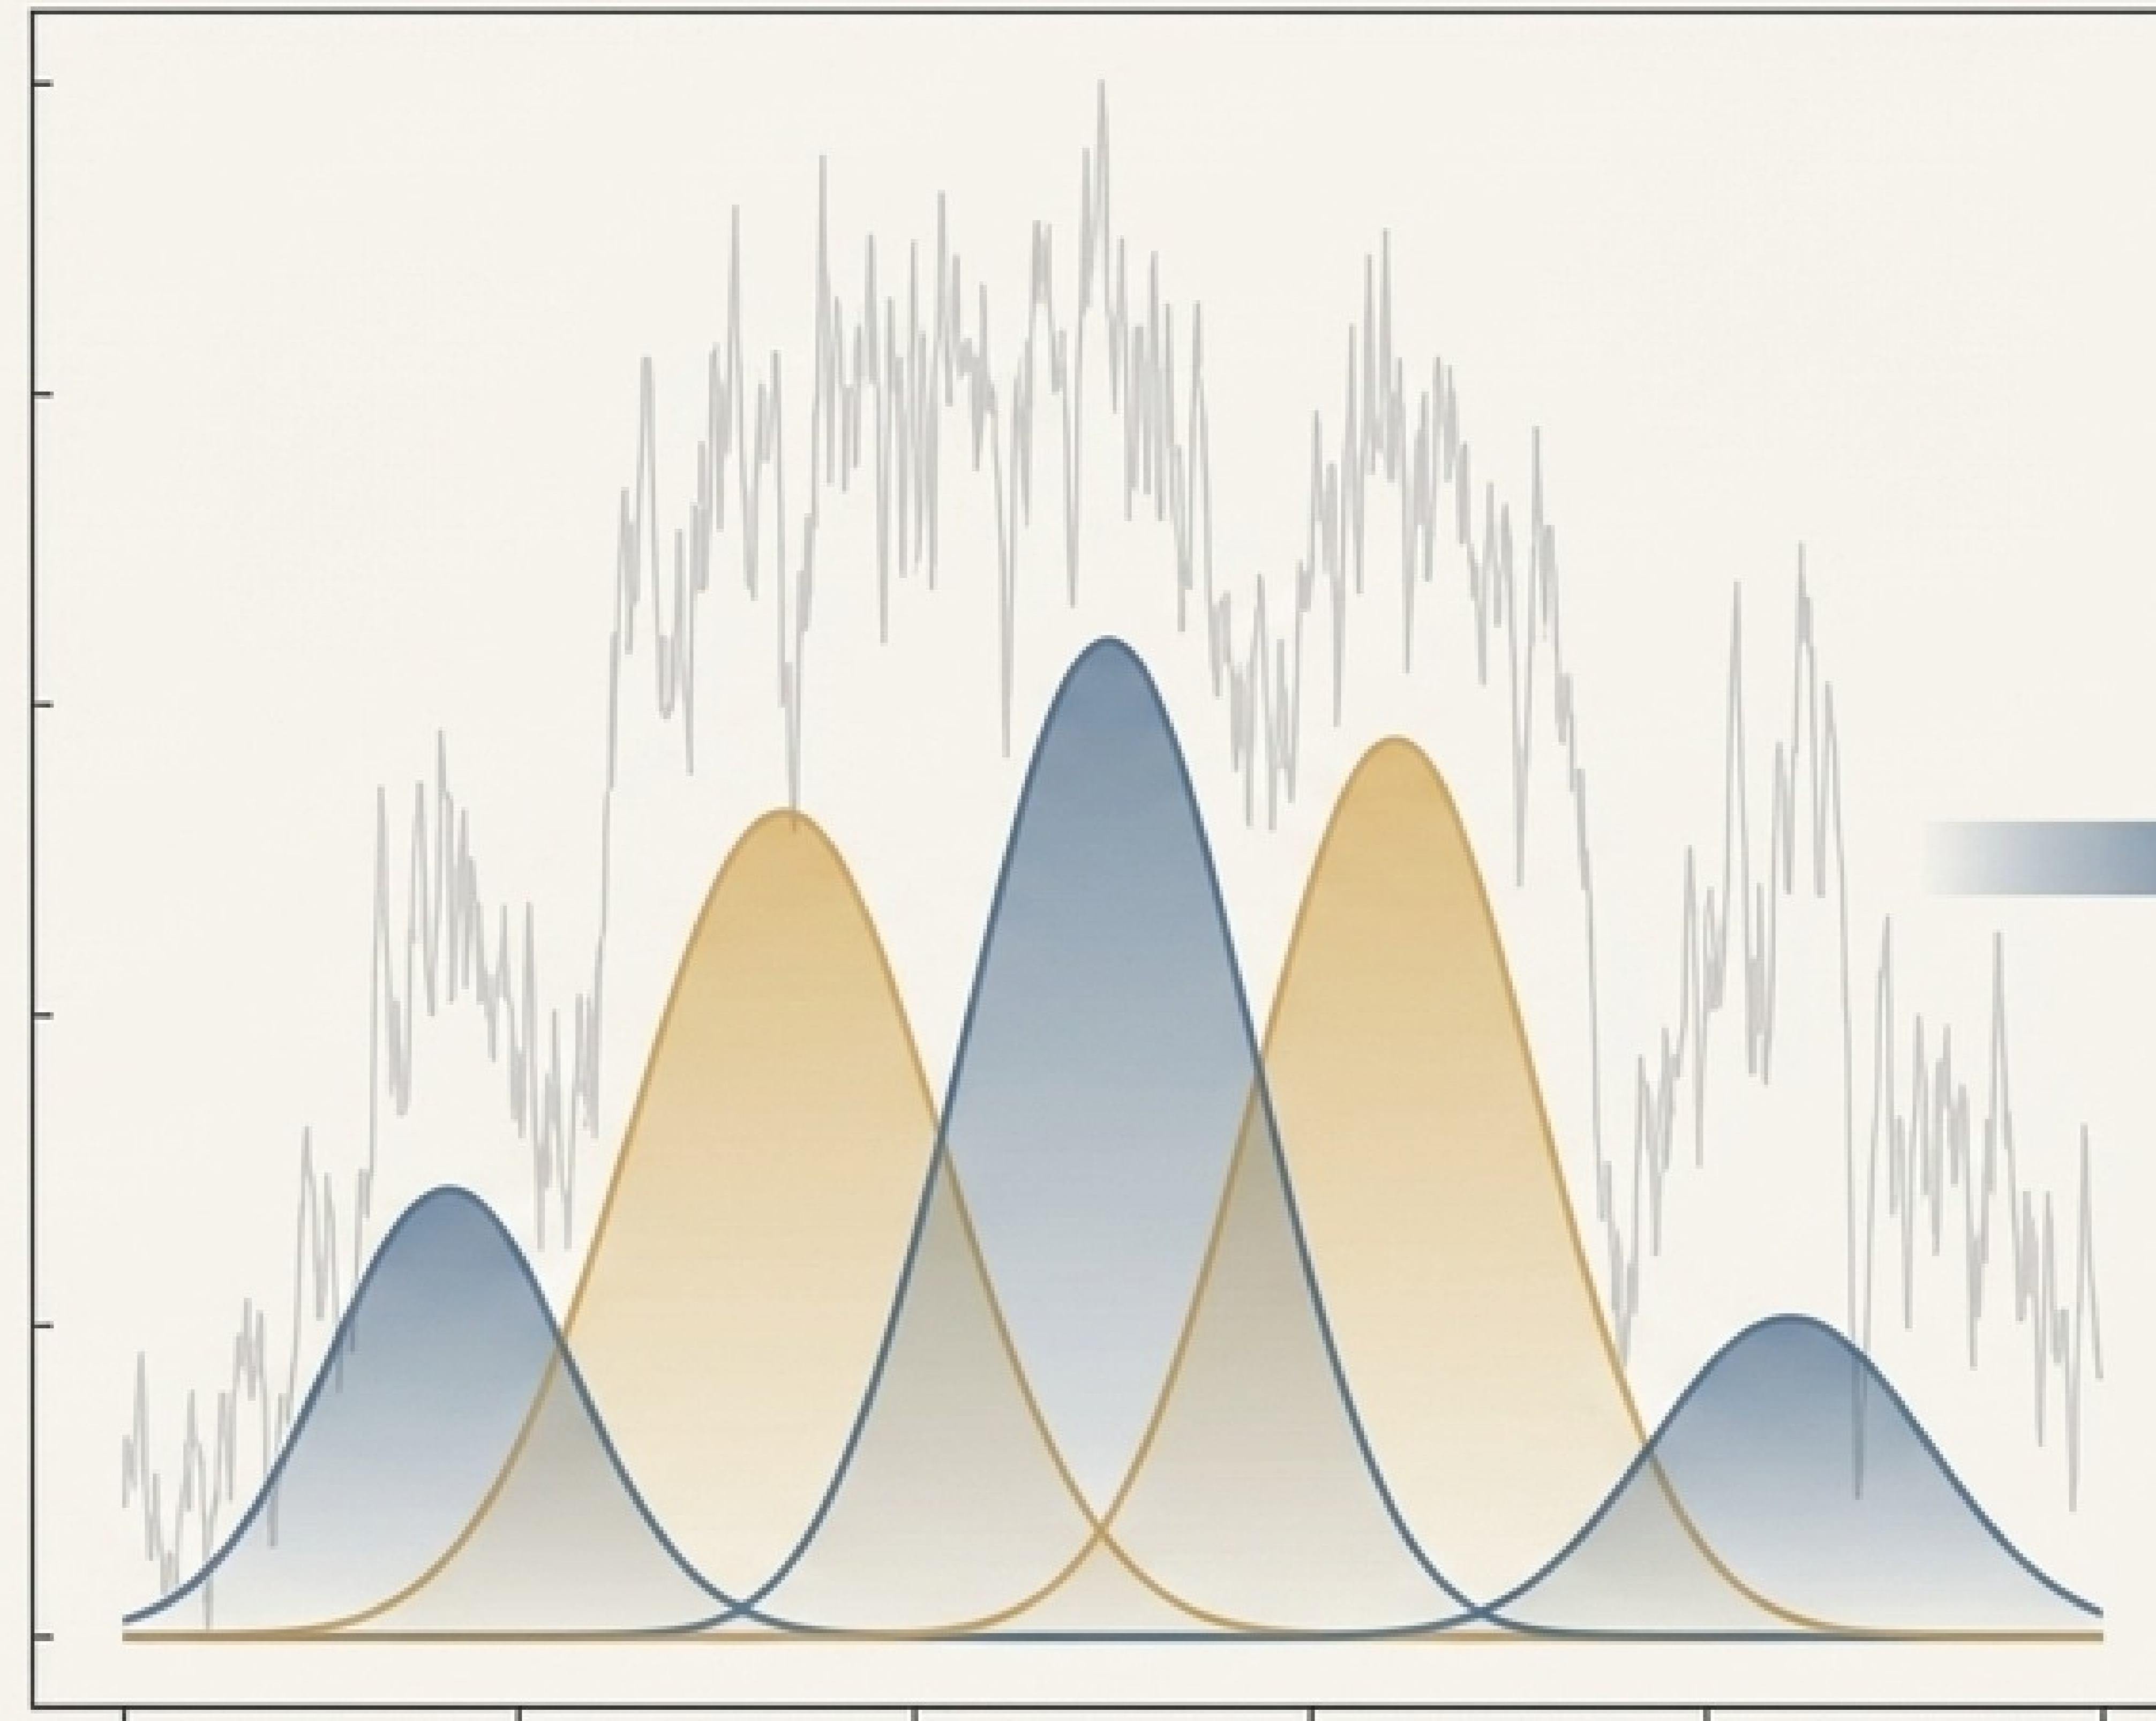
Decoding the Earth: Hyperspectral Imaging

The Challenge

Hyperspectral remote sensing captures hundreds of wavelengths per pixel, easily overwhelming traditional CNNs.

Physics-Informed Spectral Modelling (PhISM)

Disentangles observations using continuous basis functions.



Skew Normal Parameters

Instead of raw data, the AI learns to predict just 4 physical parameters per spectral component:

1. Mean (μ)
2. Standard Deviation (σ)
3. Skew (α)
4. Scale (s)

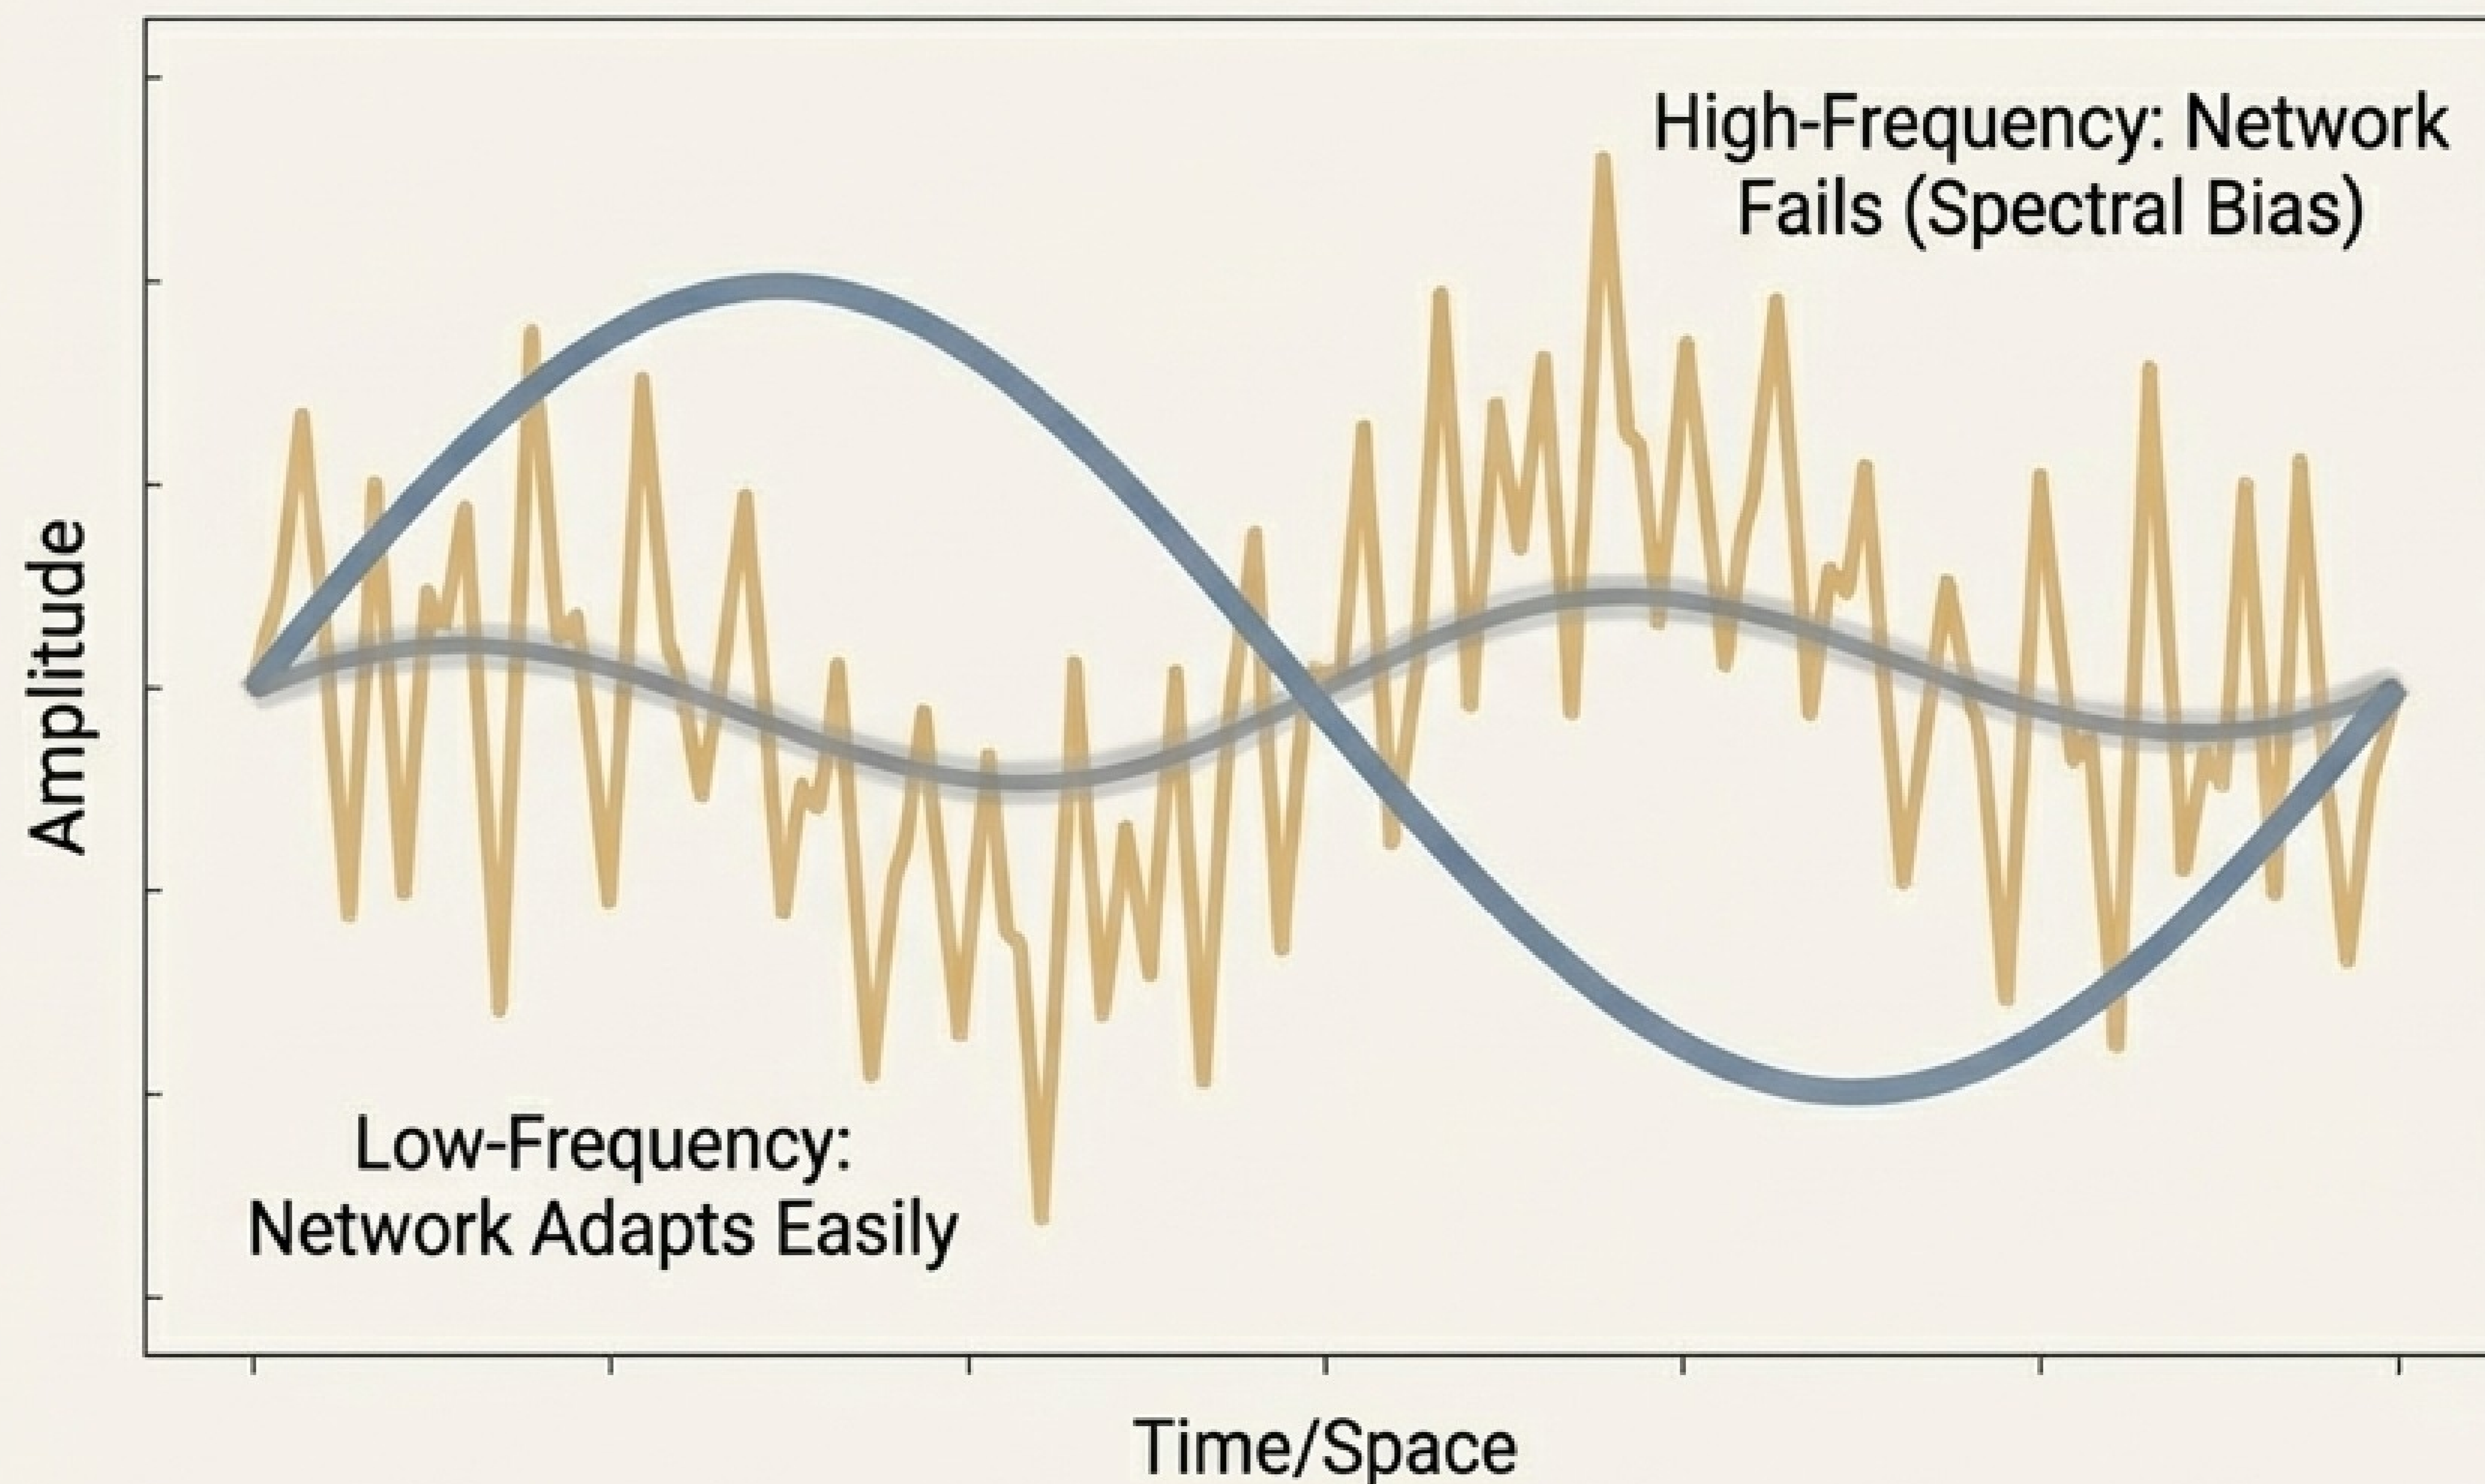
The AI Resolution Challenge

Implicit Neural Representations (INRs)

Parameterising continuous signals via Multi-Layer Perceptrons (MLPs) allows for resolution-independent modelling of 3D shapes, audio, and images.

The Spectral Bias

Standard MLPs inherently struggle to learn high-frequency content, leading to blurry reconstructions, slow convergence, and loss of fine spatial detail.



Accelerating INRs with Meta-Learning

The Bottleneck

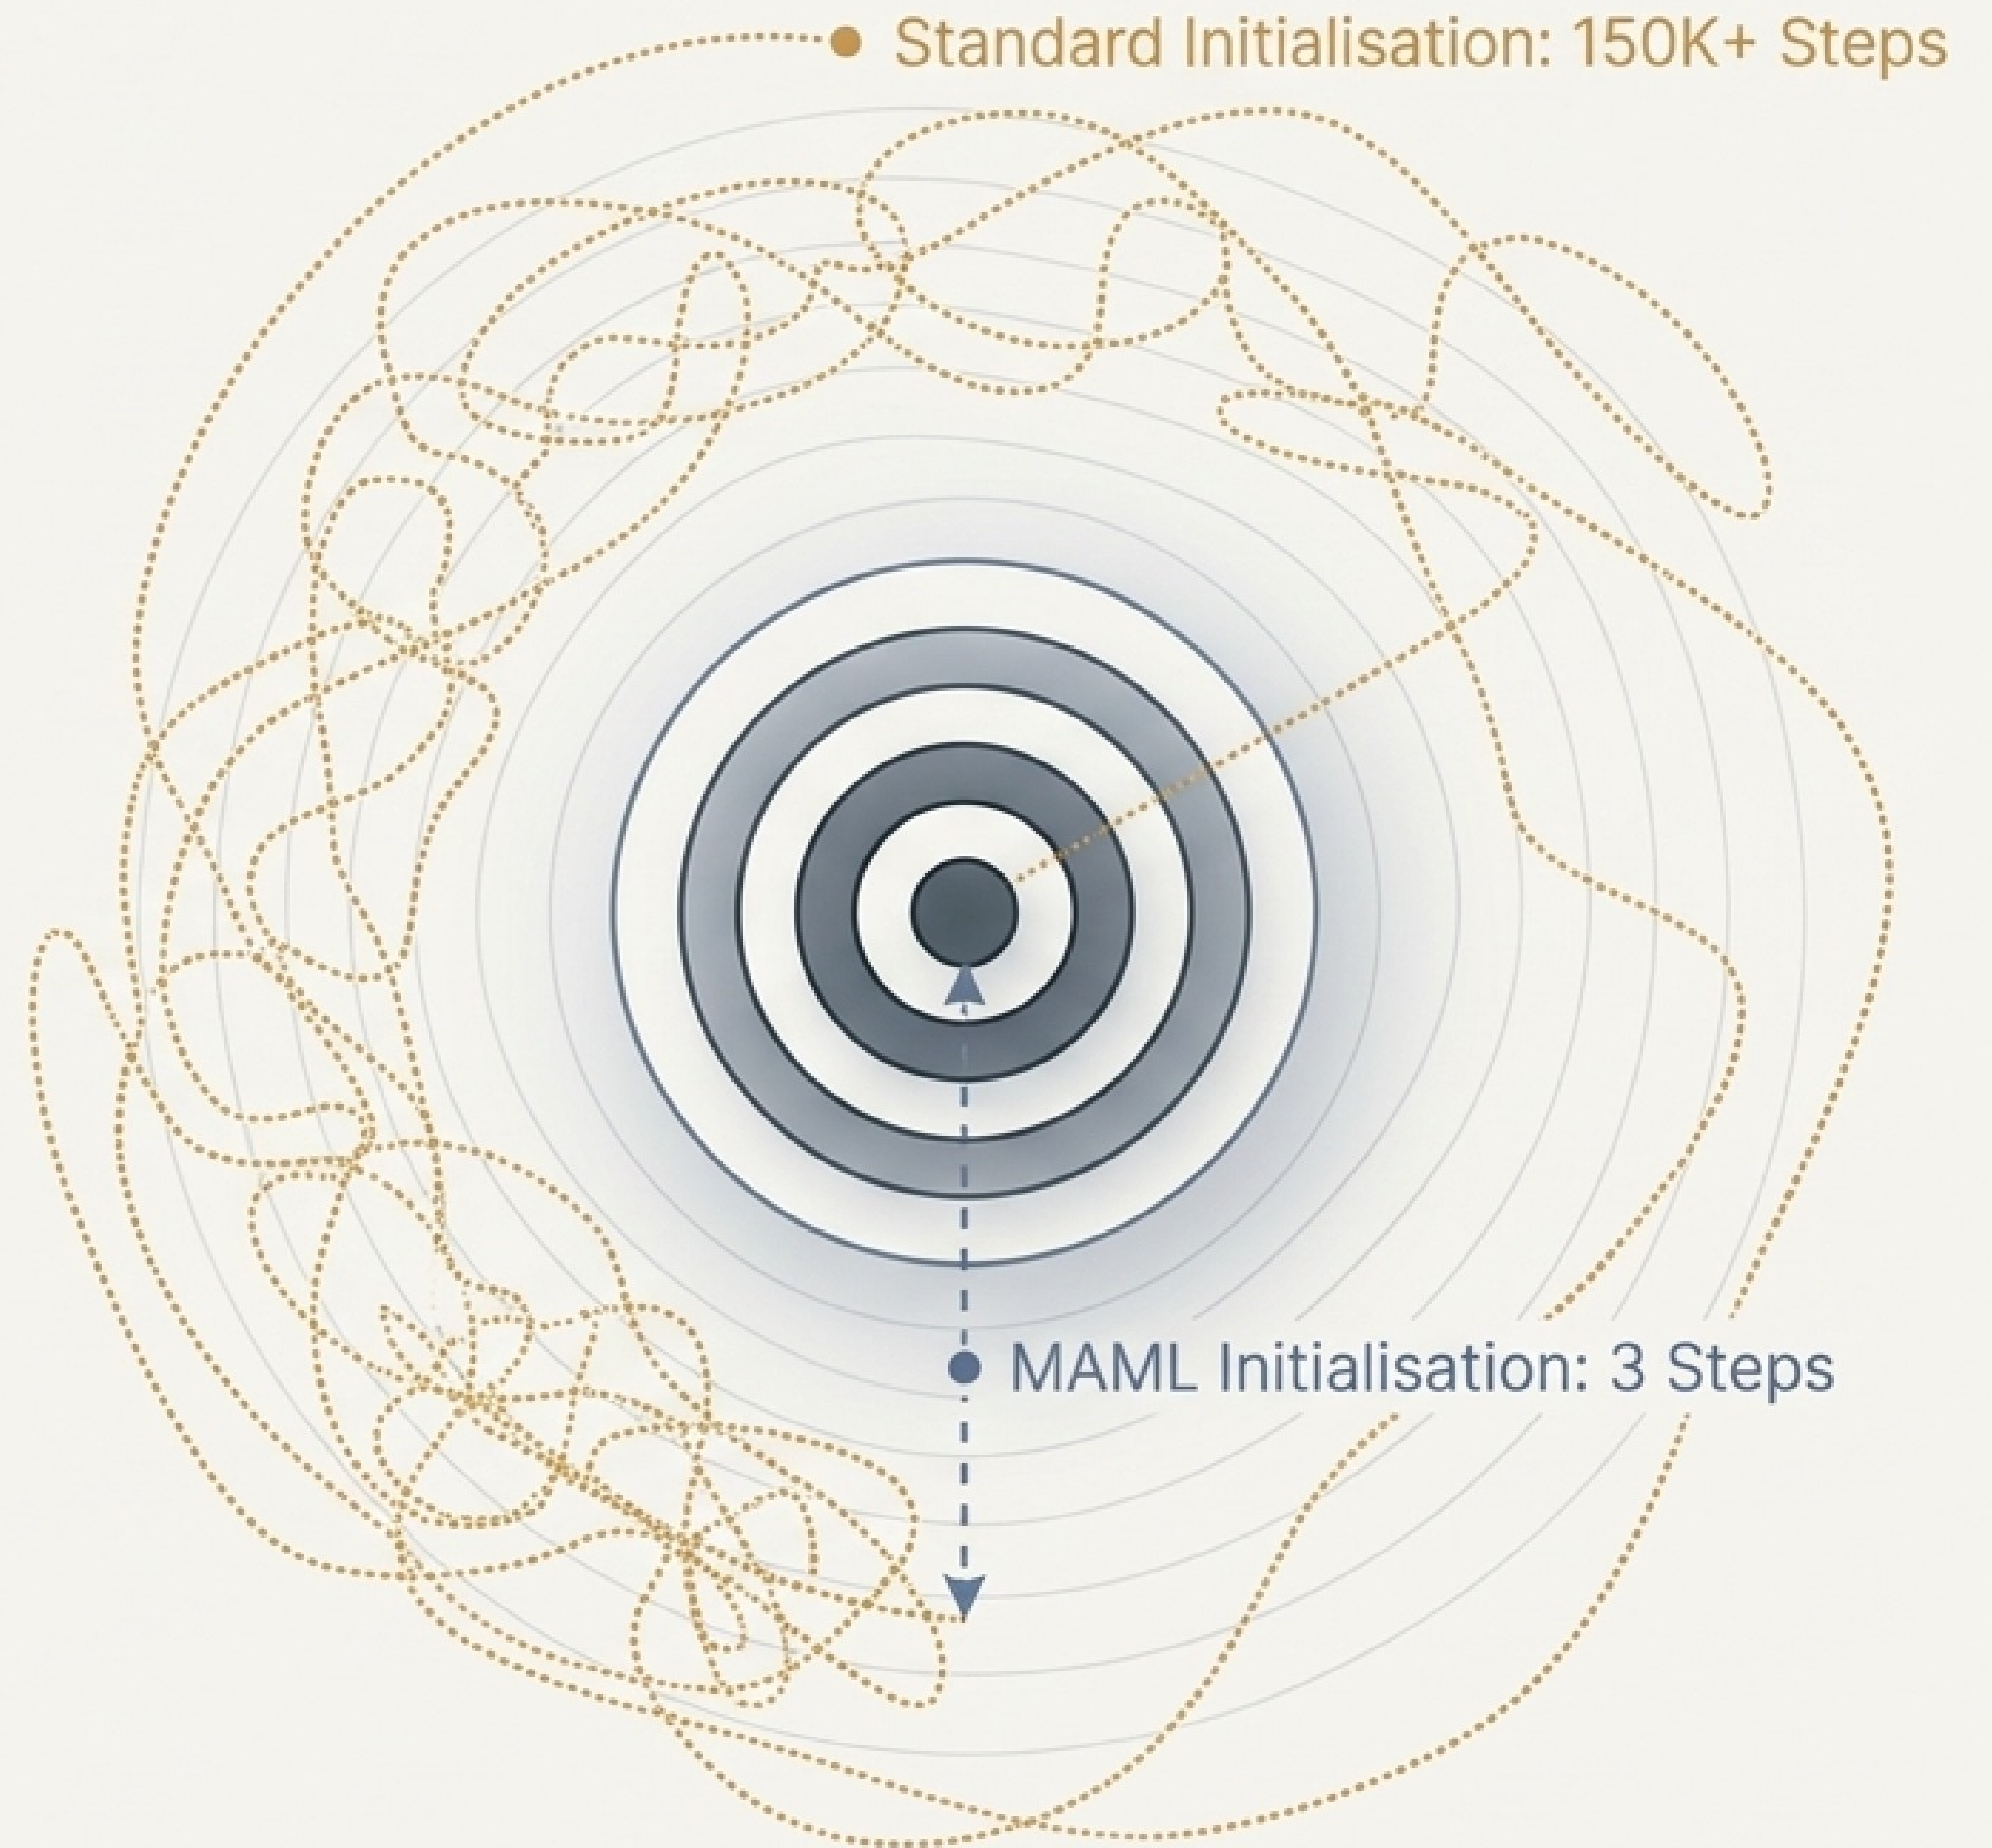
Fitting a single Neural Field to a high-resolution signal can take over 150,000 gradient steps.

Model-Agnostic Meta-Learning (MAML)

By defining an optimisation process over the network's initialisation, the AI learns a universal prior for a class of signals (e.g., human faces or 3D environments).

The Result

Adaptation to entirely novel signals requires only a handful of gradient updates, drastically reducing computational load.



The Spectral Matrix



Shaping Sound

- FFT isolates transient spikes
- Dynamic EQ targets exact frequencies
- Preserves harmonic integrity



Sculpting Light

- 2D Fourier transforms decouple magnitude from phase
- High/Low layer separation
- Seamless spatial retouching



Training AI

- INRs model continuous signals
- Physics-informed basis functions beat spectral bias
- Meta-learning accelerates spatial generation

Across acoustics, optics, and artificial intelligence, mastering the invisible dimension of frequencies is the **ultimate key to manipulating reality.**