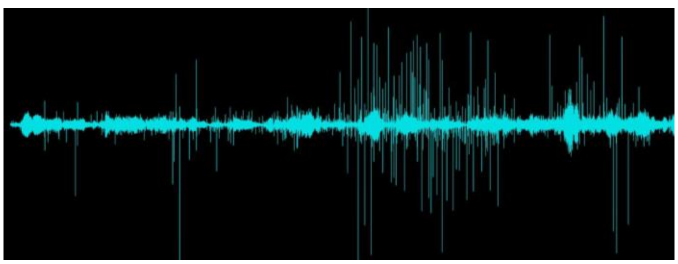

WAV ANALYZER

User Manual

ABSTRACT

This manual describes WAV Analyzer, a comprehensive desktop application for underwater acoustic analysis, calibration, and signal processing. Designed for researchers, engineers, and technicians, WAV Analyzer integrates a full suite of tools—waveform inspection, FFT‑based spectral analysis, spectrogram segmentation, hydrophone calibration curve generation, SPL transmit analysis, ambient‑noise and duty‑cycle measurements, cepstral echo detection, and more–into a single, interface. This manual guides you through each module's controls, underlying mathematics, and practical workflows, offering best‑practice tips and real-world use cases to help you harness the full power of WAV Analyzer for marine acoustics, sonar research, data management, and environmental monitoring.

Contents

1. Overview

Waveform Analyzer is a desktop application for the detailed analysis of underwater acoustic recordings. It is suitable for both technical and non-technical users including marine scientists, acousticians, data scientists, engineers, and field technicians. It supports tools for waveform manipulation, acoustic measurement, classification, spectrogram display, SPL estimation, and more.

- Key Features:

WAV File Analysis Tools

- Open and analyze WAV files
- Select channels, zoom, pan, and play audio
- Overlay plots of filtered and raw data
- Annotate intervals for measurement

Hydrophone Calibration

- Import known SPL reference curves
- Record tone sequence into WAV file
- Compute sensitivity (dB re 1 V/μPa) at each tone
- Plot calibration results (dark/light theme)
- Save and export hydrophone curves to database and CSV

Measurement Tools

- Short-Time RMS (sliding window energy)
- Peak detection with prominence filtering
- Crest Factor computation for transient detection
- SNR Estimator using selected signal/noise regions
- Octave-band analysis using ISO 266 center bands

- Duty Cycle measurement for impulsive signals

SPL & TVR Analysis

- Calculate SPL using hydrophone sensitivity or transducer TVR
- Single channel transmit analysis
- Multi-channel projection for array systems
- Overlay SPL curves across files

Graphing & Visualization

- Real-time or static spectrogram (FFT-based)
- Overlay detected features (e.g., peaks, LFM, HFM)
- Export plots in dark or light mode (JPG)
- Calibration curves, SPL curves, waveform plots

Detection & Classification

- Detect LFM/HFM sweeps and compute slope/duration
- Cepstrum-based echo spacing detection
- Event clustering with unsupervised k-means
- Machine learning classification

Batch Project Tools

- Matrix view of all files × tools with Run/View toggles
- Summary statistics tab for global overview
- Auto-analyze feature to run measurements in bulk

Database Tools

- View/edit all results in SQLite database
- Filter by frequency, date, tool, or file
- Annotate results (e.g., "clipped", "whale")
- Export to Excel or CSV

Real-World Use Cases

Field Hydrophone Calibration

Calibrate deployed hydrophones in situ using known reference signals (e.g. piston phone or calibration projector). Record a swept-tone or discrete tones at a fixed distance, then extract and level-shift the sensitivity curve so your sensors' response is traceable to a standard.

Transducer Calibrations

Characterize underwater projectors or transducers in a tank or open-water range by playing test signals (chirps, tone bursts). Use LFM/HFM analysis to map sensitivity vs. frequency and beam-pattern mapping by rotation.

Measuring SPL of a Transducer or Other Signal

Convert hydrophone voltages to SPL (dB re 1 µPa) with your calibration curve and distance correction. Generate SPL-vs-frequency CSVs, log into the database, and compare different transmitters under identical conditions.

Characterizing Signals in PAM Data

Scan long Passive Acoustic Monitoring recordings for target events (e.g. animal calls, vessel noise). Extract spectrogram clips, run duty-cycle or octave-band analyses, and log frequency, RMS, and ambient levels over time.

Detecting and Extracting LFM Sonar Pings

Automatically find broadband chirp pulses by thresholding the envelope or spectrogram. Slice each ping into a file, compute its instantaneous sweep law for matched-filter design, and log time-of-arrival and energy metrics.

Comparing Clipped vs. De-clipped Waveforms

Assess ADC clipping by comparing raw clipped recordings against restored waveforms (via interpolation). Run identical analyses (RMS, FFT, SPL) on both to quantify distortion and validate de-clipping.

Detection of Signals or Calls in Passive Acoustic Data

Use peak-prominence or event-clustering to group clicks, whistles, or vocalizations.

Compute inter-click intervals, frequency stability, and duty cycles for each call train, then export labeled clips for classification.

Ambient Noise Studies

Over long recordings, compute VRMS or time series with the Short-Time RMS tool. Apply octave-band or multitaper PSD to derive L_{90}/L_{10} statistics, spectral slopes, and exceedance levels for environmental assessments.

Filtering Data Prior to Processing

Apply notch (50/60 Hz), anti-aliasing, or FFT-based filters to remove wind noise, wave slap, or mains hum. Cleaned data improves detection, spectral estimation, and matched-filter performance.

Bulk Data Analysis

Run batch LFM/HFM or interval analyses on hundreds of WAVs overnight. Automatically log measurements (frequency, RMS, bandwidth, max voltage) into the database and generate summary CSVs or reports.

Preparing Data for Long-Term Storage

Downsample bit depth and sampling rate to reduce file size, then archive with metadata (calibration curve, analysis parameters). Keep both raw and processed versions for future audits. Automatic tools like the “Recommend Sample Rate” popup tool allow the user to make data driven decisions.

Running Machine-Learning Models for Classification & Detection

Use the Spectrogram tab to slice and label high-amplitude events (clicks, pings). Save ML clips into `ml/` folders for training neural-network classifiers.

Echo Detection & Measurement

In Depth Sounder or Active Sonar tools, band-pass around the ping frequency, compute the envelope to find direct and echo peaks, and convert time delays to range (depth). Log echo amplitude, target strength, and bandwidth.

Beam-Pattern Mapping

Measure how a transducer’s sensitivity varies with angle by rotating it in an array and logging SPL vs. bearing. Use multiple runs and overlay SPL curves to build a 3D beam pattern.

Source-Level Estimation

Derive absolute radiated sound power of transducers or marine mammals by combining SPL, distance, and beam-pattern data. Apply distance corrections and integrate over solid angle for total source level.

Underwater Noise Mapping
Deploy multiple hydrophones, grid their positions, interpolate ambient-noise levels spatially, and produce "noise heatmaps" for habitat assessment or regulatory reporting.

Acoustic Localization & Tracking
Use time-difference-of-arrival (TDOA) between hydrophones to triangulate moving vessels or tagged animals in real time. Combine with instantaneous frequency tracking for robust identification.

Reverberation & Multipath Analysis
Quantify echo delays and amplitudes in complex environments (harbors, canyons) to model reverberation time (RT_{60}) and multipath structure, informing sonar performance predictions.

Impulse-Response / Transfer-Function Measurement
Play a broadband MLS or swept sine, record the response, then deconvolve to obtain the system or channel impulse response. Analyze reflections, resonances, and time-domain artifacts.

Sonar Bathymetry & Echo-Sounding
Automate depth measurements by pinging and detecting bottom echoes along survey lines. Log depth vs. position for seabed mapping, integrating with GPS when needed.

Sensor Drift & Health Monitoring
Periodically run calibration routines on deployed hydrophones to track sensitivity drift over months or years. Log calibration curves in the database for trend analysis.

Statistical Noise Characterization
Compute long-term noise statistics (PDFs, exceedance levels like L_{90}/L_{10}) to characterize ambient conditions and inform impact assessments or sensor deployment strategies.

Filter & Algorithm Validation
Benchmark digital filters (anti-alias, notch, band-pass) and detection algorithms using synthetic test signals or known standards. Compare output metrics to expected values.

Automated Bulk-Processing Pipelines
Schedule batch jobs to process entire archives: extract features, generate spectrograms, log measurements, and refresh dashboards without manual intervention.

Machine-Learning Dataset Generation
Slice, label, and export thousands of event clips (clicks, calls, pings) for deep-learning training. Use SpanSelector and Save ML Clip to build balanced, high-quality datasets.

2. User Interface Overview

The application is divided into several sections accessible via tabs and toolbars.

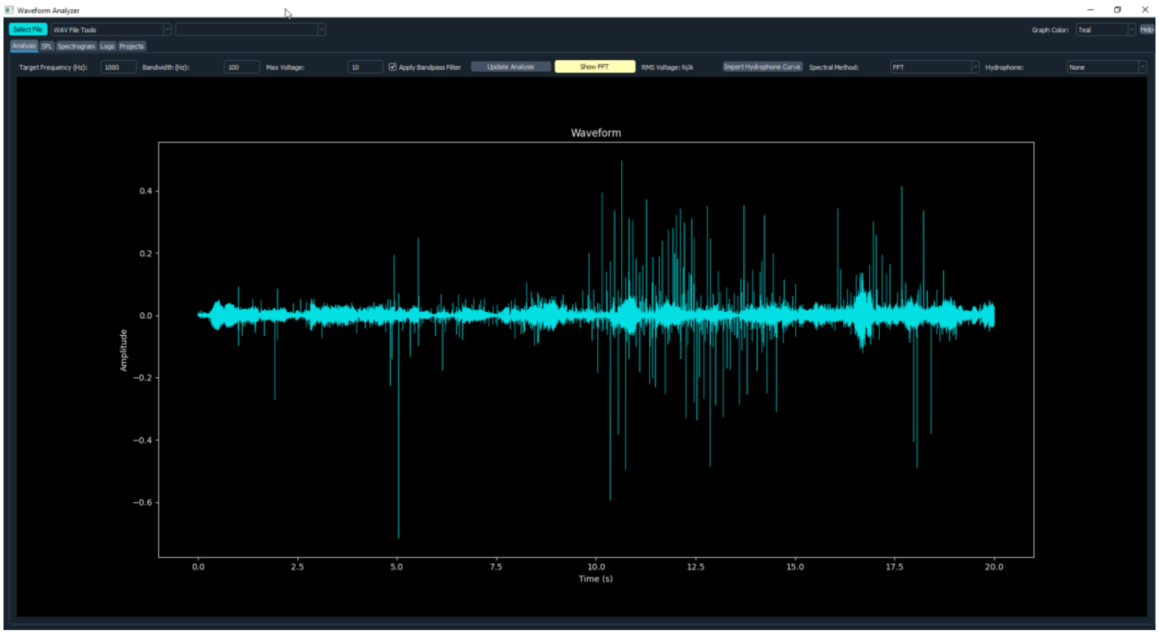

Figure 1: Main Window

Top Toolbar Functions:

- Select File – Load a WAV file
- Category Dropdown – Choose tool category
  - WAV File Tools
  - Measurement Tools
  - Detection & Classification
  - Database Tools

- Tool Dropdown – Launch selected tool
- Graph Color – Set the plot color
  - Available colours:

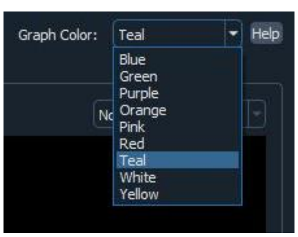

Figure 2: Graph Colours

- Help – Open help content

Main Tabs

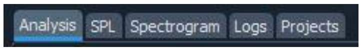

Figure 3: Main Tabs

- Analysis – Primary waveform plot and controls
- SPL – Displays SPL results from voltage measurements
- Spectrogram – Time-frequency display
- Logs – Tabular view of saved results
- Projects – File × Method matrix tracker

FFT Mode Panel

Toggled on via the FFT toolbar button. This panel appears beneath the waveform and lets you inspect spectral content in a moving window.

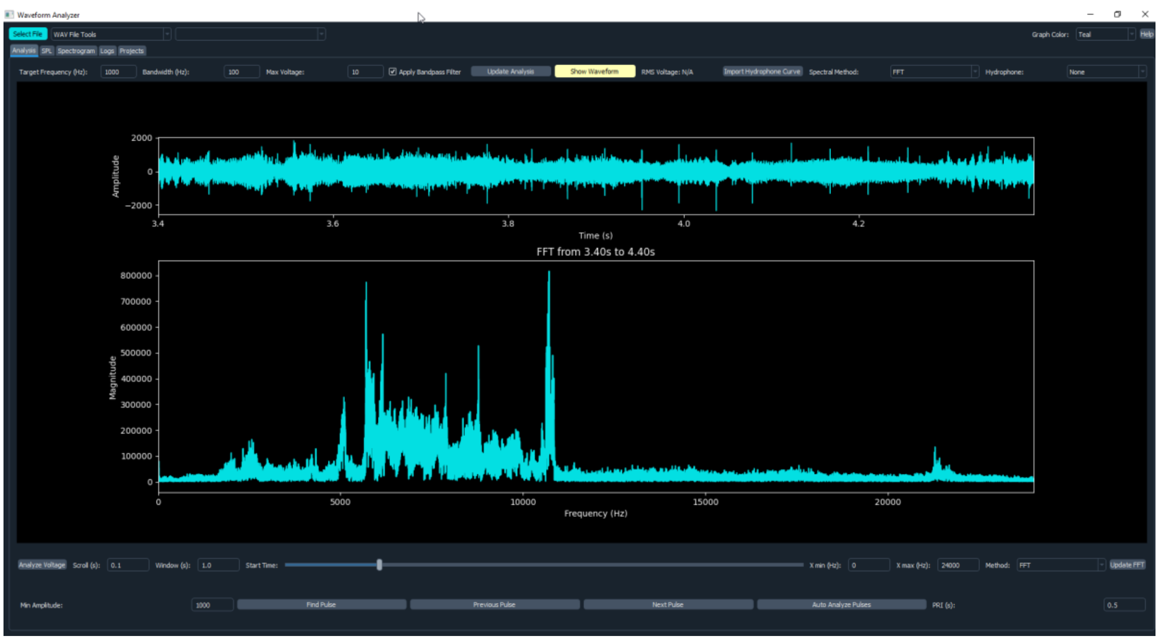

Figure 4: FFT Window

Control	Description
Window Length (s)	Duration of each FFT slice. Affects $\Delta f = \frac{1}{window}$.
Start Time Slider	Move the FFT window start time across the file; updates spectrum live.
Threshold (for Pulse Mode)	Minimum peak amplitude in the window to continue LFM/HFM analysis; stops sweep when below threshold.

Control	Description
Min Voltage Entry	Maximum DAQ voltage (e.g. 10 Vp) used to convert integer samples into volts before FFT.
Discard 0 Hz	When enabled, any window whose dominant frequency = 0 Hz is skipped.
Spectral Plot	Frequency on X, magnitude (dB) on Y; axes and grid adapt to theme.

How to use FFT mode

1. Load a WAV and switch to **FFT** mode.

2. Set **Window Length** and **Threshold (min amplitude)** as needed.

3. Drag the **Start Time** slider to your region of interest, the spectral plot updates in real time. Optionally, you can set the min amplitude, eg 1000, and click “find pulse”.
    1. Once you have found the section you want to analyze, you can adjust your window to ensure you measure the correct part of your wav file.
        1. The window begins at the beginning of the amplitude plot, see Figure 5: Amplitude Plot, and ends at the end of the amplitude plot.

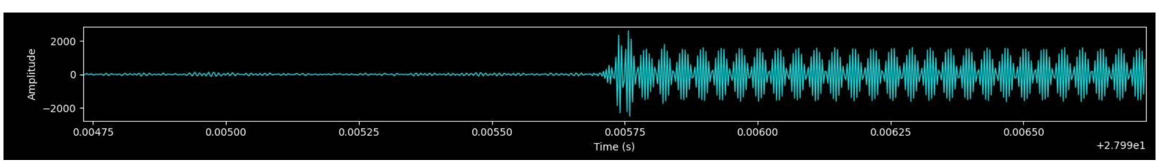

Figure 5: Amplitude Plot

        2. Align the start of the amplitude plot with the start of the signal you want to measure.
            1. Use the “scroll steps” to align your window with the start of the signal. Enter the time you want to scroll and then use the arrow keys to move forwards and backwards in time.

Figure 6: Scroll Steps Input

3. Click the “Analyze Voltage” button.

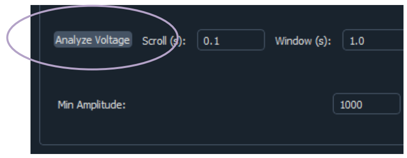

Figure 7: Analyze Voltage Button

How to measure signal length

1. Using the raw Amplitude plot, click on the start of the signal.
2. Then click on the end of the signal.
3. A popup will display showing the start time in seconds, stop time in seconds, Δt in seconds, distance in meters, and the dominant frequency in Hz.

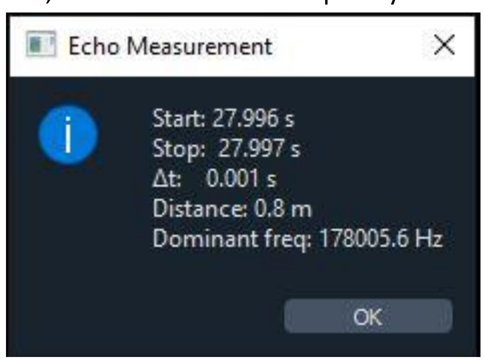

Figure 8: FFT Measurement Popup

Filter & Spectral Analysis Methods

Below are six core techniques, each with the mathematical background, real-world use cases, and a glossary of variables.

---

1. FFT-Based Filtering

How It Works

DFT:

$$X[k] = \sum_{n=0}^{N-1} x[n]e^{-j\frac{2\pi}{N}kn}$$

Apply Mask:

$$Y[k] = H[k]X[k]$$

Inverse DFT:

$$y[n] = \frac{1}{N}\sum_{k=0}^{N-1} Y[k]e^{j\frac{2\pi}{N}kn}$$

Use Cases

Notch Filtering: Remove narrowband interference (e.g. 50/60 Hz mains hum).

Band-pass Selection: Isolate a tight frequency band, e.g. only 1 kHz test tone.

Variable Glossary

- $x[n]$: input signal samples.
- N: number of DFT points (FFT size).
- $X[k]$: spectrum of x.
- $H[k]$: frequency-domain filter mask (0/1 or shaped).
- $Y[k]$: filtered spectrum.
- $y[n]$: time-domain filtered output.

---

2. Hilbert Transform & Envelope Extraction

How It Works

Compute analytic signal:

$$z(t) = x(t) + jH\{x(t)\},$$

where H{·} is the Hilbert transform.

Instantaneous envelope:

$$A(t) = |z(t)| = \sqrt{x(t)^2 + (H\{x(t)\})^2}$$

Use Cases

Duty-Cycle Analysis: Measure time spent above an amplitude threshold using the envelope.

Onset Detection: Find pulse start by a sharp rise in $A(t)$.

Phase-Based Modulation: Extract instantaneous phase for FM or phase-shift analysis.

Variable Glossary

$x(t)$: continuous signal.

$H\{x(t)\}$: Hilbert transform.

$z(t)$: complex analytic signal.

$A(t)$: envelope (instantaneous amplitude).

3. Welch's Method for Power Spectral Density

How It Works

Segment & Window each length-LLL block with w[n]w[n]w[n]:

$$P_m[k] = \frac{1}{LU}|DFT\{w[n]x_m[n]\}|^2, U = \frac{1}{L}\sum_{n=0}^{L-1} w[n]^2$$

Average over MMM segments:

$$P_{Welch}[k] = \frac{1}{M}\sum_{m=1}^{M} P_m[k].$$

Use Cases

Ambient Noise Spectra: Smooth, low-variance PSD estimates over long recordings.

Comparative Studies: Compare noise floors across sites or times of day.

Quality Assurance: Verify that system noise remains within specifications.

Variable Glossary

$x[n]$: input samples.

$x_m[n]$: mth overlapping segment.

$w[n]$: window function.

L: segment length.

M: number of segments.

$P_m[k]$: periodogram of segment m.

$P_{welch}[k]$: averaged PSD.

U: window power normalization.

4. Multitaper Spectral Estimation

How It Works

Use K orthogonal tapers $V_k[n]$ (Slepian sequences).

Compute each tapered spectrum and average:

$$S_{MT}(f) = \frac{1}{K}\sum_{k=1}^{K} |DFT\{x[n]V_k[n]\}|^2$$

Use Cases

Short Recordings: Accurate PSD when data length is limited.

Low-Variance Spectra: Resolve narrowband features in noisy data.

Precision Calibration: Characterize fine spectral features of transducers.

Variable Glossary

$x[n]$: input signal.

$V_k[n]$: kth taper.

K: number of tapers.

$S_{MT}(f)$: multitaper PSD estimate.

5. Instantaneous Frequency

How It Works

From $A(t)e^{j\emptyset(t)}$, extract phase $\emptyset(t) = \arg(z(t))$.

Differentiate:

$$f_{inst}(t) = \frac{1}{2\pi}\frac{d\emptyset(t)}{dt}.$$

Use Cases

Chirp Tracking: Follow LFM/HFM sweep laws continuously.

FM Signal Analysis: Decode frequency-modulated communications.

Non-Stationary Signals: Reveal rapid frequency changes in transient events.

Variable Glossary

$z(t)$: analytic signal.

$\emptyset(t)$: instantaneous phase.

$f_{inst}(t)$: instantaneous frequency (Hz).

6. Modulation Frequency Analysis

How It Works

Envelope $\boldsymbol{A(t)}$ via Hilbert (Sec 2).

Spectrum of $A(t)$(FFT or Welch) to reveal modulation rates.

Use Cases

Pulse-Train Rates: Detect repetition rates of pulses or clicks.

Vibration Analysis: Identify mechanical modulation frequencies.

Bioacoustic Rhythms: Quantify calling rates or burst patterns in marine life.

Variable Glossary

$A(t)$: envelope (instantaneous amplitude).

Modulation frequencies: spectral peaks of $A(t)$.

Selection Guide

Method	**Time Res vs. Freq Res**	**Variance**	**Best For**
FFT Filter	low Δt, high Δf	N/A	Sharp frequency-domain filtering
Hilbert	–	–	Envelope & phase extraction
Welch	moderate Δt/Δf	low	Ambient-noise PSD
Multitaper	high Δt/Δf	very low	Short-record, high-precision spectra
Instantaneous Freq	continuous t–f	can be noisy	Chirp and FM analysis
Modulation Spectrum	Δf of envelope	depends	Repetition-rate detection

3. WAV File Tools

• Trim File

Removes a specified duration from the beginning and/or end of the WAV file. Useful for removing silence or irrelevant sections.

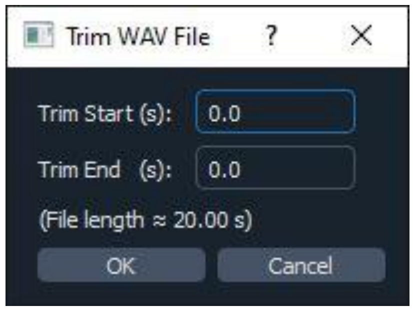

Figure 9: Trim File Popup

• Denoise File

Applies a Wiener filter to reduce background noise:

$$x(n) = \frac{S_x(f)}{S_x(f) + S_n(f)} * y(n)$$

Figure 10: Denoise Dialog Popup

The Wiener filter is a classic, statistically optimal denoising method that adapts itself to your signal's local SNR. Below are its primary **use cases** and the **effects** you can expect on your logged measurements:

Use Cases

Low-SNR Pulse Detection

When your hydrophone pulse is buried in broadband ambient noise (e.g. shipping or wind-driven sea noise), Wiener filtering can boost pulse clarity, helping you detect and time-stamp low-level echoes.

Ambient Noise Analysis

To track ambient VRMS or spectrum over long recordings, Wiener filtering removes random fluctuations—yielding smoother VRMS/time curves and more stable dominant-frequency estimates.

LFM Sweep Processing
During linear-frequency sweeps, background noise can vary across the band. The Wiener filter's frequency-dependent gain preserves sweep energy while suppressing noise floor, improving your RMS and FFT-based frequency picks.

Effects on Measurement Values

Metric	Typical Wiener Effect
RMS Voltage	Slightly **reduced** (noise floor removed), but **more repeatable**. Over-smoothing is rare at moderate filter lengths.
Peak Voltage	May be **smoothed down** if a sharp spike is mistaken for noise—choose window lengths cautiously to avoid clipping true peaks.
Dominant Frequency	**Stabilized** across noisy segments; small shifts (<1 Hz) possible if filter over-attenuates narrowband components.
Bandwidth	**Narrower** apparent bandwidth (noise tails suppressed), yielding cleaner bandwidth logging but potentially missing faint sidelobes.
Time-of-Arrival / Duty Cycle	**Sharper** transition detection when noise is high, but filter delay may introduce a small, constant time offset—account for this if you need sub-millisecond accuracy.

• Highpass Filter

Applies a Butterworth highpass filter to remove low-frequency noise:

$$H(s) = \frac{s^n}{s^n + a_1 s^{n-1} + \cdots . + a_n}$$

Figure 11: Highpass Filter Dialog Popup

• Lowpass Filter

Applies a Butterworth lowpass filter to remove high-frequency content.

Figure 12: Lowpass Filter Dialog Popup

• Anti-Aliasing Filter

Prepares data for down sampling. Filters frequencies above 0.45× new sample rate.

$$f_c = 0.45 * f_{new\ sample\ rate}$$

Figure 13: Anti-Aliasing Filter Popup

An anti-aliasing filter is a low-pass filter applied before downsampling or in hardware to prevent high-frequency components from folding into lower frequencies. Here's how it fits into your hydrophone workflow:

Use Cases

Downsampling WAV Files
Before reducing the sample rate (e.g. from 192 kHz to 48 kHz), you must remove content above the new Nyquist frequency (½ fs) to avoid spectral aliasing.

Real-Time Streaming & Recording
Many ADCs include an analog anti-aliasing stage; when you add a digital low-pass filter in software, you ensure that no out-of-band noise re-enters your analysis.

Band-Restricted Measurements
If you know your hydrophone is only valid up to, say, 50 kHz, applying a sharp digital anti-aliasing filter at that cutoff prevents any energy above 50 kHz from contaminating your VRMS or FFT results.

Effects on Measurement Values

Metric	**Typical Anti-Aliasing Effect**

Metric	Typical Anti-Aliasing Effect
RMS Voltage	**Reduced** slightly by removing high-frequency noise; generally results in a **truer** measurement of energy within the passband.
Peak Voltage	**Unchanged** for pulses well below cutoff; very fast transients near cutoff may be **smoothed**, lowering the measured peak.
Dominant Frequency	**Accurate** within passband; prevents spurious high-frequency folding that could produce false peaks below Nyquist.
Bandwidth	**Limited** to the filter's cutoff; out-of-band sidelobes are removed, yielding a clean, defined bandwidth.
Time-Domain Waveform	Introduces a small **group delay** (filter latency) and **edge smoothing**—account for this delay if measuring precise time-of-arrival.
Duty Cycle & Leq	**More stable** over time as high-frequency noise is suppressed, but very short high-frequency spikes may be lost, slightly lowering duty-cycle percentages.

• Recommend Sample Rate

Estimates the minimal sample rate required based on spectral content.

$$f_s = 2 * f_{\max detected}\ (Nyquist\ criterion)$$

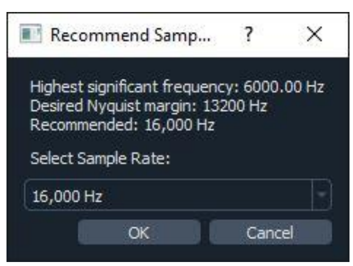

Figure 14: Recommend Sample Rate Popup

This popup analyzes your loaded WAV file, finds the highest "significant" frequency content, and then offers a suitable standard sampling rate to use when downsampling. Below are its primary **use cases** and the **effects** that resampling will have on your logged measurement values.

Use Cases

File Size & Storage

You have a very long recording or high-rate WAV (e.g. 192 kHz) and want to reduce disk space or memory usage.

Hardware / Format Compatibility

Your downstream tools or DAC/ADC only support certain standard rates (44.1 kHz, 48 kHz, etc.).

Nyquist Margin

Your content only goes up to, say, 20 kHz—there's no need to keep 96 kHz sampling. You can safely drop to 44.1 kHz or 48 kHz.

Processing Speed

FFTs, filtering, and plotting all run faster at lower sample rates. Downsampling can substantially accelerate batch analyses.

Effects on Measurement Values

Measurement	Effect of Resampling at a Lower Rate
RMS Voltage & SPL	• **Minor amplitude changes** due to interpolation/filtering (usually < 0.1 dB)

• Remove DC Offset

Subtracts the mean from the signal:

$$x_{new}(n) = x(n) - (\frac{1}{n}\sum_{n=1}^{N} x(n)$$

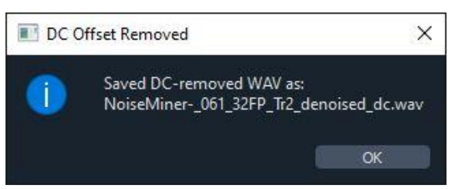

Figure 15: DC Offset Popup

Subtracting the DC offset (the mean value) from your hydrophone time-series ensures your waveform is centered around zero. Here's when you'd apply it and what it does to your logged metrics:

Use Cases

Sensor & Cable Bias

Hydrophones and pre-amplifiers often introduce a constant voltage bias (DC) due to electronics or cabling. Removing that bias prevents artificial shifts in your voltage measurements.

Long-Duration Recordings

Over time, temperature changes or power-supply drift can cause the recording baseline to wander. A DC-offset correction (on each file or in short windows) stabilizes your baseline.

FFT & Spectral Analysis

Any nonzero mean injects a large amplitude at 0 Hz in your FFT. Removing DC ensures your spectral plots and dominant-frequency detection aren't dominated by that spike.

Effects on Measurement Values

Metric	Effect of Removing DC Offset
RMS Voltage	• **Slightly reduced**, because you're subtracting out constant energy.

• Normalize File

Scales amplitude so that the max sample reaches full scale:

$$x_{norm}(n) = \frac{x(n)}{\max |\, x(n)}$$

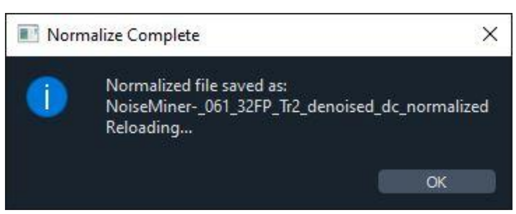

Figure 16: Normalize Dialog Popup

• Downsample Bit Depth

Purpose:

Convert floating-point audio to 16-bit or 8-bit integer format to reduce file size or simulate hardware.

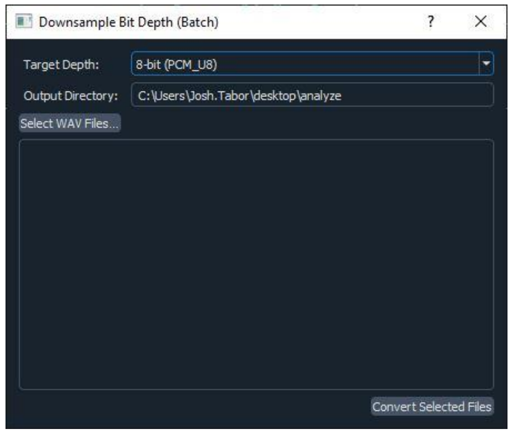

Figure 17: Down Sample Bit Depth Popup

Reducing a file's bit depth changes its digital resolution—trading dynamic range and precision for smaller file sizes or hardware compatibility. Below are when you'd use this feature and how it affects your logged measurements.

Use Cases

File Size & Storage

Long deployments or continuous recordings at 32-bit can consume huge disk space. Dropping to 24- or 16-bit slashes file size by 25–50%.

Hardware/Software Compatibility

Some audio interfaces, editors, or legacy systems only support 16-bit or 24-bit WAVs.

Downstream Processing

If you know your measurement's dynamic range never exceeds, say, 12 bits (≈72 dB), you can safely resample to 16-bit without losing acoustic information.

Archival Standards

Many long-term archives standardize on 24-bit at 48 kHz. Bit-depth conversion lets you meet those requirements.

Effects on Measurement Values

Metric	**Effect of Reducing Bit Depth**
RMS Voltage & SPL	• **Minor quantization noise** raises the noise floor by ~½ LSB RMS.

How to Use

1. Choose *Downsample Bit Depth*.
2. Select desired bit depth (e.g., 16-bit).
3. The file is re-quantized:

$$x_{int}(n) = round\ (x(n) *\ 2^{(B-1)}$$

Where B is the number of bits.

4. Measurement Tools

• Ambient Noise

Estimates background RMS level in a specified quiet window.

Equation:

$$VRMS = \sqrt{\frac{1}{N}\sum_{n=1}^{N} x(n)^2}$$

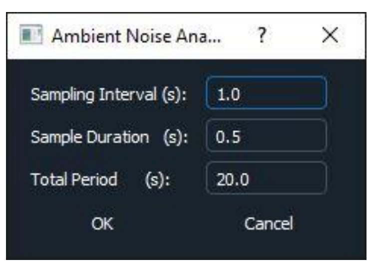

Figure 18: Ambient Noise Popup

The Ambient Noise tool samples your file at regular intervals to chart how background acoustic energy (VRMS) and dominant frequency evolve over time. This helps characterize environmental noise levels, identify temporal patterns, and compare sites or conditions.

Use Cases

Long-Term Monitoring
Track diurnal or tidal cycles in background noise (e.g. shipping, wind-driven waves) over hours or days.

Site Comparison
Compare noise profiles between multiple recording locations or deployment periods (e.g. before vs. after construction).

Equipment Diagnostics
Verify that your hydrophone and pre-amp noise floor remains stable over time, or detect electronic noise anomalies.

Masking & Thresholding
Use the noise curves to set adaptive detection thresholds for pulse analysis or event detection (e.g. marine mammal calls).

Effects on Measurement Values

Metric	How Ambient-Noise Analysis Affects It
VRMS (AC)	• **Smoothed over window**: short-duration spikes are averaged out, yielding a stable background level per sample period.

• Peak Prominences Analysis

Detects peaks that rise above surrounding noise using prominence-based detection.

Useful for detecting impulses or echolocation clicks.

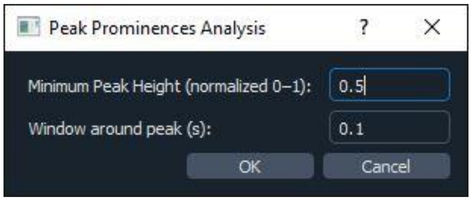

Figure 19: Peak Prominences Analysis Popup

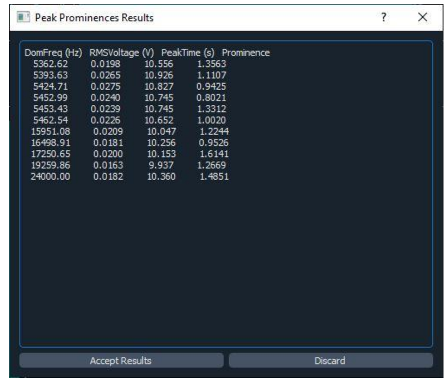

Figure 20: Peak Prominences Analysis Popup Results

• Interval Analysis

Measures RMS and dominant frequency within a user-specified time interval.

Dominant frequency calculated from FFT magnitude:

$$f_{dom} = \arg\max |\, FFT\ (x(t))$$

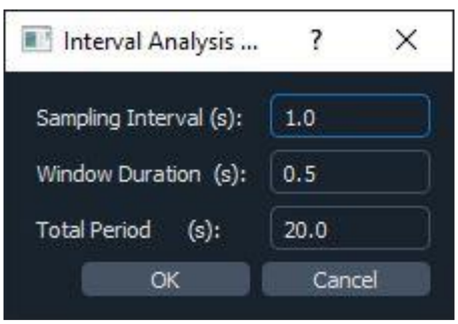

Figure 21: Interval Analysis Popup Single Files

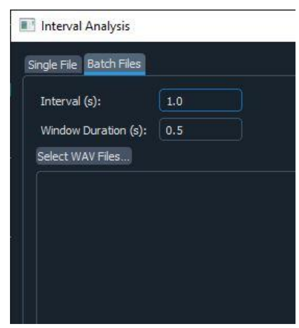

Figure 22: Interval Analysis Popup Batch Files

• Depth Sounder Analysis

Identifies pings and their echoes to calculate time-of-flight:

ToF = t_echo - t_ping

Range = $\frac{c*ToF}{2}$,

where c = speed of sound in water (~1500 m/s)

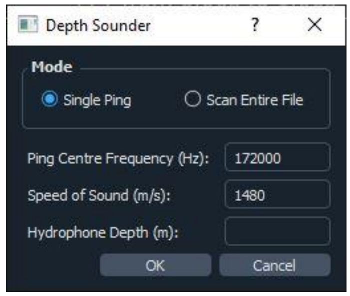

Figure 23: Depth Sounder Analysis Popup

The “Peak Prominence” popup lets you set how “tall” a peak must be—relative to its surroundings—to be counted. This is essential whenever you need to discriminate true acoustic pulses from spurious bumps or noise.

Use Cases

Pulse Detection in Noisy Environments
– When background noise has random fluctuations, raising the prominence threshold prevents false-positive pulses from being logged.

Separating Closely-Spaced Pulses
– If two pulses overlap, a lower prominence lets you detect smaller secondary peaks; raising it merges them into a single event.

Echo & Reflection Analysis
– To capture only the strongest echo returns (and ignore minor reverberations), increase the prominence so only high-energy reflections are counted.

Automated Threshold Tuning
– Use the popup interactively: slide the prominence value while viewing your waveform and spectrogram to find the “sweet spot” between over- and under-detection.

Effects on Measurement Values

Metric	**Effect of Changing Peak Prominence**
Number of	Higher prominence → **fewer** peaks (only the most distinct); lower → **more**,

Metric	Effect of Changing Peak Prominence
Detected Peaks	including small ripples.
Peak Voltage	Only peaks exceeding the prominence barrier are measured. Smaller local maxima are **ignored**, so minimum recorded peak voltage **increases**.
Pulse Timing (Time-of-Arrival)	Shifting prominence can **merge** or **split** adjacent peaks, altering the logged arrival times for pulses or echoes.
Duty Cycle	If duty cycle is based on time above threshold around each peak, changing which peaks are counted will **raise or lower** the calculated duty percentage.
Dominant Frequency	When you use prominence in spectral-peak selection (e.g. in FFT mode), a higher setting filters out low-amplitude spectral lobes, yielding a **cleaner**, but possibly **less detailed**, frequency trace.
Repeatability & Uncertainty	Interactive tuning lets you quantify how many peaks "drop out" as you vary prominence—use that range to estimate your detection uncertainty.

• Slope De-Clipper

This tool attempts to restore clipped waveforms using linear extrapolation and slope continuity across clipped regions.

It is suitable for both soundcard clipping (±1.0) and preamp clipping (not reaching ±1.0).

Process:
1. Detect constant-valued regions (flat tops).
2. Estimate signal slope before and after the clip.
3. Reconstruct clipped region using a linear or spline fit.

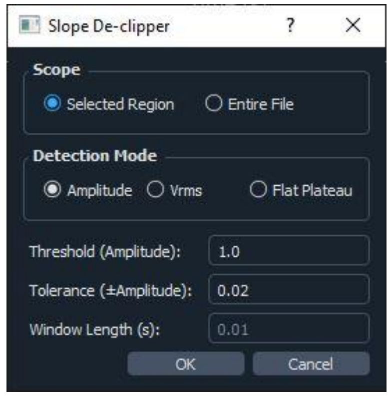

Figure 24: Slope De-clipper Popup

The Slope Declipper popup identifies "flat-topped" clipped portions of your waveform—where the hydrophone or ADC saturated—and attempts to reconstruct the true signal shape by interpolating or smoothing across those regions.

Use Cases

Overdriven Pulses
When a loud ping exceeds your DAQ's maximum input and produces a flat plateau, the slope declipper can restore the rising/falling edges for more accurate peak and RMS measurements.

Overdriven Hydrophone Preamp
When a loud ping exceeds your hydrophone preamp maximum input and produces a flat plateau, the slope declipper can restore the rising/falling edges for more accurate peak and RMS measurements.

Transient Spike Recovery
Short, high-amplitude transient events (e.g. motor knocks, bubble collapse) can clip; declipping helps you recover their true amplitudes and durations.

Calibration Validation
If you accidentally record your calibration tone at too high a level, declipping lets you salvage the run instead of re-measuring from scratch.

Batch Cleanup
Apply declipping automatically across dozens of files before downstream analysis, ensuring consistent pulse shapes and avoiding manual re-recording.

Effects on Measurement Values

Metric	Impact of Slope Declipper
Peak Voltage	**Increased** relative to clipped plateau (recovers true peak height), but may be **underestimated** if interpolation window is too short.
RMS Voltage & SPL	**Raised** slightly because the recovered portions add back energy; improves accuracy when clipping was extensive.
Dominant Frequency	**Unaffected** for steady tones; for swept or broadband signals, declipped edges can sharpen FFT magnitude but do not alter frequency content.
Time-of-Arrival	**More precise**, since the true rising-edge is reconstructed; peak detection "zero-crossings" shift back to their real positions.
Duty Cycle	**Longer** "on" times if clipping had flattened plateaus; declipping restores the original envelope so duty-cycle percent increases toward the true value.
Waveform Shape	Introduces a **smooth interpolation**—watch for overshoot or ringing if aggressive polynomial fitting is used.

• Find Peaks

Identifies and logs time/frequency of signal peaks.

Options:
- Peak height threshold
- Minimum peak distance
- Prominence

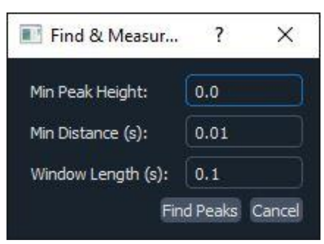

Figure 25: Find & Measure Peaks

The Find Peaks popup lets you configure how peaks are identified in your time-domain waveform. You can set parameters like **minimum height**, **minimum distance**, and **prominence** so that only the acoustic events you care about are detected.

Use Cases

Single-Pulse Extraction

To locate the main transmit pulse in a recordings of marine transducers.

Echo & Reflection Analysis

To automatically detect multiple echo returns in a sonar ping, separating each reflection.

Click / Click-Train Detection

For bioacoustics, to pick out cetacean click trains or snapping-shrimp snaps against a noisy background.

Burst & Sweep Characterization

To identify individual peaks in LFM sweeps or pulsed bursts for individual frequency segment analysis.

Effects on Measurement Values

Measurement	Impact of Peak-Finding Parameters
Number of Peaks	↑ with lower height/distance thresholds; ↓ with higher thresholds or larger minimum-distance settings.
Peak Voltage (Height)	Only peaks above your set **minimum height** or **prominence** are reported, so raising the threshold **raises** the minimum logged amplitude.
Pulse Timing	`peak_index / fs` gives time-of-arrival. A larger **minimum distance** avoids detecting jittery neighbors, yielding cleaner timing but possible missed closely spaced reflections.
Inter-Peak Interval	Controlled by **minimum distance**: setting it too low can record spurious “sub-peaks” as separate events, artificially reducing your measured interval.
Duty Cycle	If duty cycle depends on time above threshold around each peak, changing peak selection alters the “on” durations you include—raising height or prominence typically **lowers** duty cycle.
Repeatability	Consistent parameter settings ensure that small noise fluctuations don’t shift which peaks get detected, improving measurement reproducibility.

• Short-Time RMS

Computes moving-window RMS:

V_RMS(t) = sqrt((1/N) * sum[x²(t - i)])

Use cases:
- Transient signal strength
- Whale call energy tracking

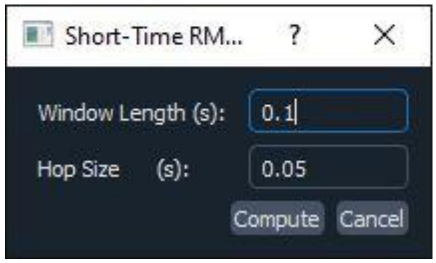

Figure 26: Short Time RMS Popup

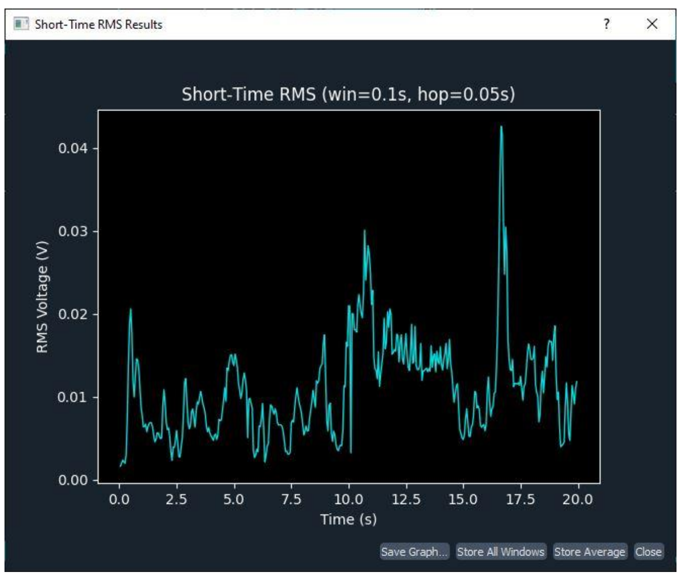

Figure 27: Short Time RMS Popup Results

This tool computes the Root-Mean-Square (RMS) level (and equivalent continuous level, Leq) in short, overlapping time windows—letting you track how acoustic energy varies over time at a fine granularity.

Use Cases

Transient Event Characterization

Identify when short bursts (e.g. vessel pass-bys, snapping shrimp snaps) cause sudden rises in background levels.

Continuous Noise Monitoring
Chart how ambient noise (wind, waves, machinery) fluctuates minute-by-minute or second-by-second.

Comparing Signal vs. Noise
Overlay the RMS envelope of a sweep or ping on top of background RMS to see SNR variations throughout a measurement.

Regulatory Compliance
Compute Leq over standard intervals (e.g. 1 s, 10 s) for environmental noise reporting or permit requirements.

Effects on Logged Measurement Values

Metric	Effect of Short-Time RMS Processing
RMS Voltage	• **Window-averaged** voltage—rapid spikes are smoothed within each window.

• Octave Band Analysis

Splits spectrum into 1-octave or 1/3-octave bands and computes RMS per band.

Filters:
- Log-spaced bandpass filters

Follows ISO 266 center frequencies.

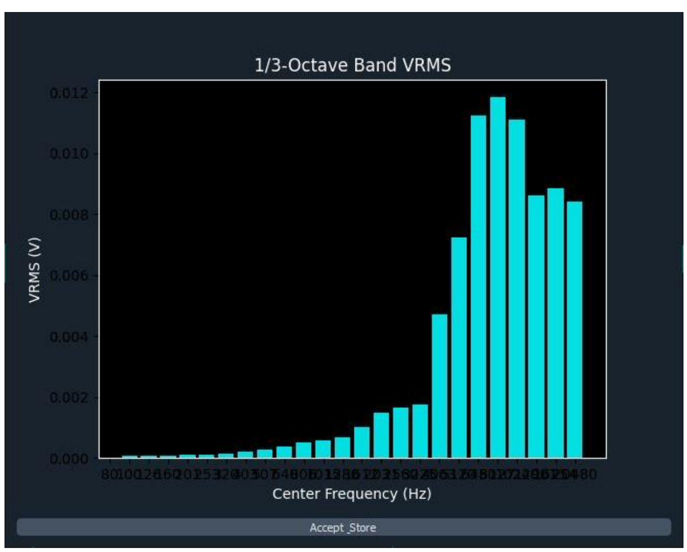

Figure 28: Octave Band Analysis Popup Results

Octave-band analysis breaks your spectrum into standardized frequency bands (e.g. 1-octave or 1/3-octave), computing the Sound Pressure Level (SPL) within each band. This gives you a band-limited view of acoustic energy, useful for noise metrics and system diagnostics.

Use Cases

Regulatory Noise Compliance
– Many standards (e.g. ANSI, ISO) specify limits per 1-octave or 1/3-octave band (e.g. environmental or workplace noise).

Frequency-Dependent Sensitivity
– Characterize how your hydrophone/transducer response varies by band, or verify that your system meets manufacturer specs.

Source Identification
– Different noise sources (machinery, marine life, vessels) often occupy distinct bands—octave analysis helps isolate and track them.

Signal Processing Pre-Filtering
– Before further analysis, you can monitor band levels to adaptively set filters or thresholds per band (e.g. notch out a high-energy band).

Effects on Measurement Values

Metric	Effect of Octave-Band Processing
Band RMS / SPL	• Computes RMS level in each band → SPL = $20 \cdot \log_{10}(VRMS) + K$.

• SNR Estimator

User selects noise and signal intervals.

Then computes:

$$SNR\ (dB) = 20log_{10}(\frac{V\ Signal}{V\ Noise})$$

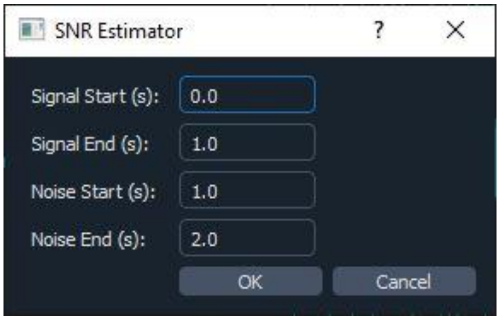

Figure 29: SNR Estimator Popup

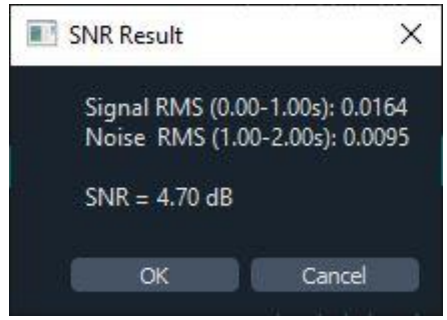

Figure 30: SNR Estimator Popup Results

The SNR (Signal-to-Noise Ratio) Estimator computes the ratio between your signal's power (or RMS level) and the background noise floor—providing a single number that summarizes measurement quality.

Use Cases

Quality Control
Confirm that each pulse or recording has sufficient SNR (e.g. >20 dB) before logging results or proceeding with further analysis.

Threshold Setting
Automatically adapt detection thresholds: raise them in low-SNR conditions to avoid false alarms, lower them when SNR is high to catch subtle events.

Equipment Comparison
Compare pre-amp or hydrophone front-end performance by measuring SNR on the same test signal across devices or gain settings.

Environmental Assessment
Track how ambient conditions (sea state, shipping traffic) affect your SNR over time or between sites.

Effects on Measurement Values

Logged Metric   How SNR Estimation Affects It

SNR Value (dB) • Computed as $20 \cdot \log_{10}(V_s / V_n)$ or $10 \cdot \log_{10}(P_s / P_n)$.

• LFM / HFM Analysis

Detects and characterizes frequency sweeps (LFM/HFM).

- Estimates start/stop frequency, duration, slope
- Uses cross-correlation with sweep kernel or Hough transform

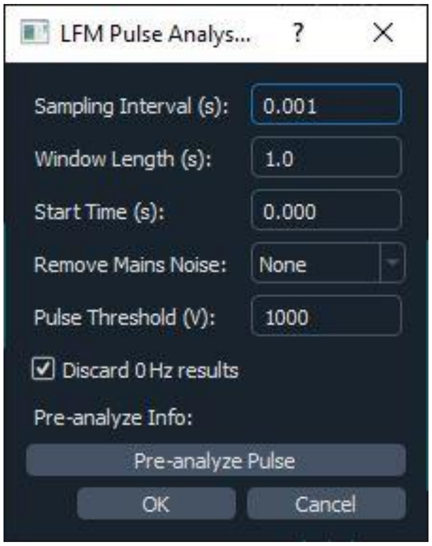

Figure 31: LFM Pulse Analysis Popup

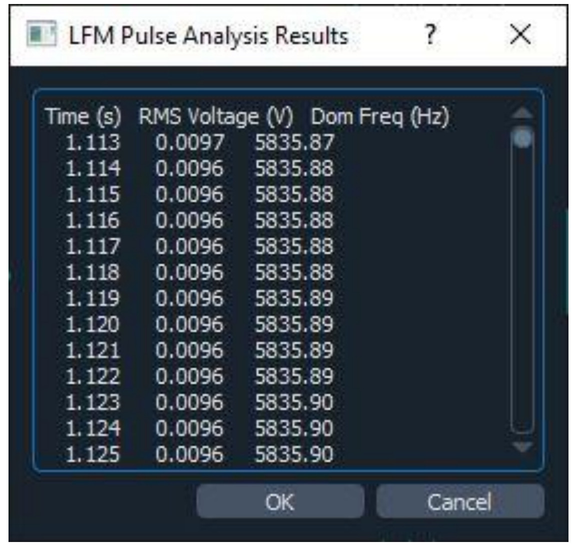

Figure 32: LFM Pulse Analysis Popup Results

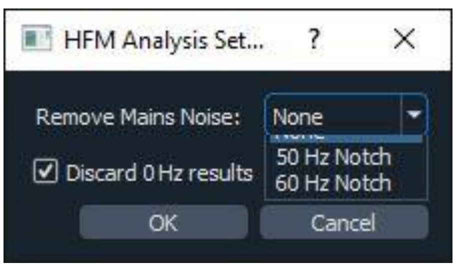

Figure 33: HFM Pulse Analysis

Figure 34: LFM Batch Analysis Popup

The LFM (Linear-Frequency-Modulated) and HFM (Hyperbolic-Frequency-Modulated) Analysis tool steps through a swept-frequency pulse in fixed windows, computing metrics like RMS voltage and instantaneous dominant frequency. It's ideal for characterizing chirp signals used in sonar, transducer calibration, and high-resolution spectral studies.

Use Cases

Chirp Transducer Calibration
Measure your transducer's sensitivity and beam pattern across a wide, continuous frequency sweep in one go.

Broadband Source Characterization
Analyze LFM/HFM pulses used in underwater imaging or seismic surveys—extracting frequency-dependent response in a single measurement.

Echo Profiling & Matched-Filter Prep
Pre-analyze chirp pulses to design matched-filter windows and thresholds for precise echo detection in cluttered environments.

Spectral Correction
Compute per-window frequency response to correct for system or propagation effects (e.g. absorption) across the entire sweep.

Effects on Logged Measurement Values

Metric	Effect in LFM/HFM Analysis
RMS Voltage (per window)	• Tracks instantaneous energy in each frequency segment—reveals any amplitude dips or resonances.
Dominant Frequency	• Pinpoints the center frequency of each window—should follow the sweep law (linear or hyperbolic).

• Multi-Frequency Analysis

Detects multiple tone components in FFT.

Steps:
1. Take FFT of target window
2. Identify N strongest peaks
3. Measure frequency, amplitude, optionally phase

Useful for tonal sonar or whistles.

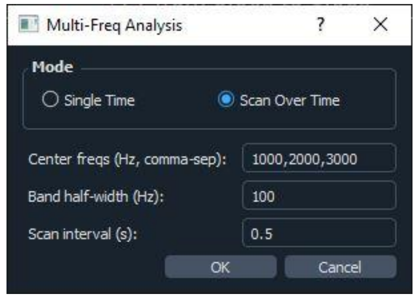

Figure 35: Multi-Frequency Analysis Popup

The Multi-Frequency Analysis tool lets you simultaneously analyze multiple discrete tones (e.g. a set of test frequencies) in a recording, extracting metrics (RMS, peak, SPL, etc.) for each frequency of interest. It's ideal when you have recordings containing several pure-tone test signals or when you want to probe specific bands in one pass.

Use Cases

Fixed-Tone Calibration
– Play a sequence of single-frequency test tones (e.g. 1 kHz, 5 kHz, 10 kHz) and log sensitivity at each frequency without reloading or repositioning.

Multi-Band Source Characterization
– Evaluate a broadband source by stepping or hopping through discrete frequencies rather than a continuous sweep.

Selective Band Monitoring
– In ambient or machinery noise surveys, track specific machine-signature frequencies (e.g. 120 Hz, 300 Hz, 1.2 kHz harmonics) over time.

Simultaneous Tone Decoding
– In underwater communications or multi-carrier sonar, measure the strength and timing of each carrier frequency in a single recording.

Effects on Measurement Values

Metric	Impact in Multi-Frequency Analysis
RMS Voltage per Tone	Measures VRMS over the tone's duration window—should match the tone's true amplitude if windowed correctly. Lower if window picks up off-tone noise.
Peak Voltage	Captures the maximum instantaneous level of each tone—good for verifying

Metric	Impact in Multi-Frequency Analysis
	crest factor and non-linearity at each frequency.
SPL (via Calibration)	Converted per-frequency using your hydrophone's TVR curve—allows plotting sensitivity vs. frequency for discrete points.
Bandwidth Logging	Minimal for pure tones (≈0 Hz bandwidth). If tones have slight modulation, the tool picks the narrowband peak only.
Time-of-Arrival	If tones start at different times, the tool logs start/end for each—helpful when tones are played in sequence or coded patterns.
Frequency Accuracy	Depends on FFT resolution vs. window length. Discrete-tone windows must be long enough (≥10 cycles) for <1 Hz accuracy.
Adjacent Tone Leakage	Window sidelobes can bleed energy from neighboring tones—use a well-tapered window (Hanning/Hamming) or notch out off-tones to improve isolation.

• SPL Transmit Analysis

Estimates SPL from RMS voltage using TVR curve.

Modes:

1. Single-channel:

$$SPL\ (f) = \ 20log_{10}(Vrms) + TVR\ (f)$$

2. Multi-channel: includes coherent summation

TVR: Transducer's dB re 1 µPa @ 1m per volt input.

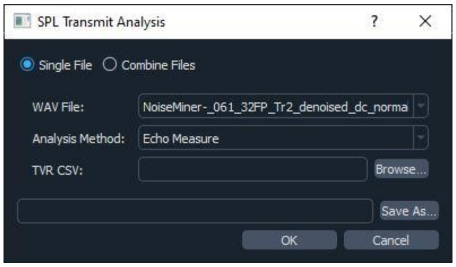

Figure 36: SPL Transmit Analysis Popup

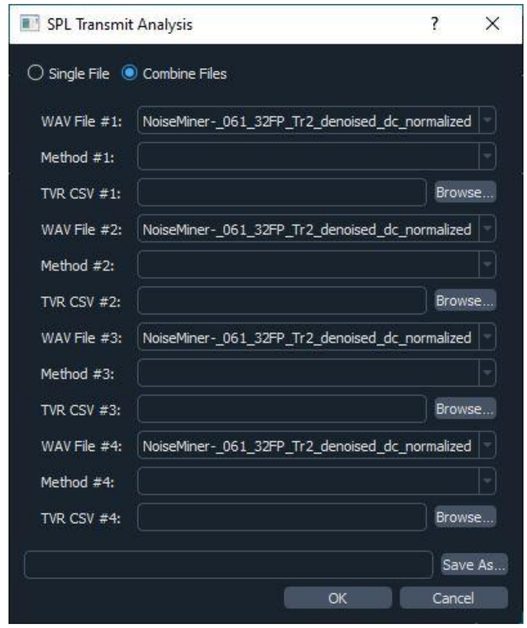

Figure 37: SPL Transmit Analysis Multi Channel Popup

The SPL Transmit Analysis tool converts your measured hydrophone voltages (VRMS or peak) into Sound Pressure Level (SPL) vs. frequency using one or more calibration curves (TVR). It supports both single-file mode and multi-file combination mode.

Use Cases

Single-Point Transmit Calibration
– You have a single WAV of a tone sweep or discrete tones and want to log SPL vs. frequency for one method.

Multi-Channel / Multi-Run Combination
– Combine up to four separate measurement runs (different files or methods) into a single SPL curve for better SNR or extended bandwidth.

Beam Pattern Characterization
– Rotate your transducer or hydrophone array, record separate files at each angle/method, then combine to produce a full-angle SPL map.

Quality Assurance & Drift Detection
– Repeat transmit tests daily or weekly and overlay multiple SPL curves to spot sensitivity drifts or anomalies.

• Hydrophone Calibration

Uses a known SPL vs frequency reference.

Steps:

1. Record WAV using calibration tones.
2. Compute RMS voltage at each tone.
3. Compute sensitivity:

$$S(f) = SPL\ Ref\ (f) - 20log_{10}VRMS\ (f)$$

Generates and stores sensitivity curve.

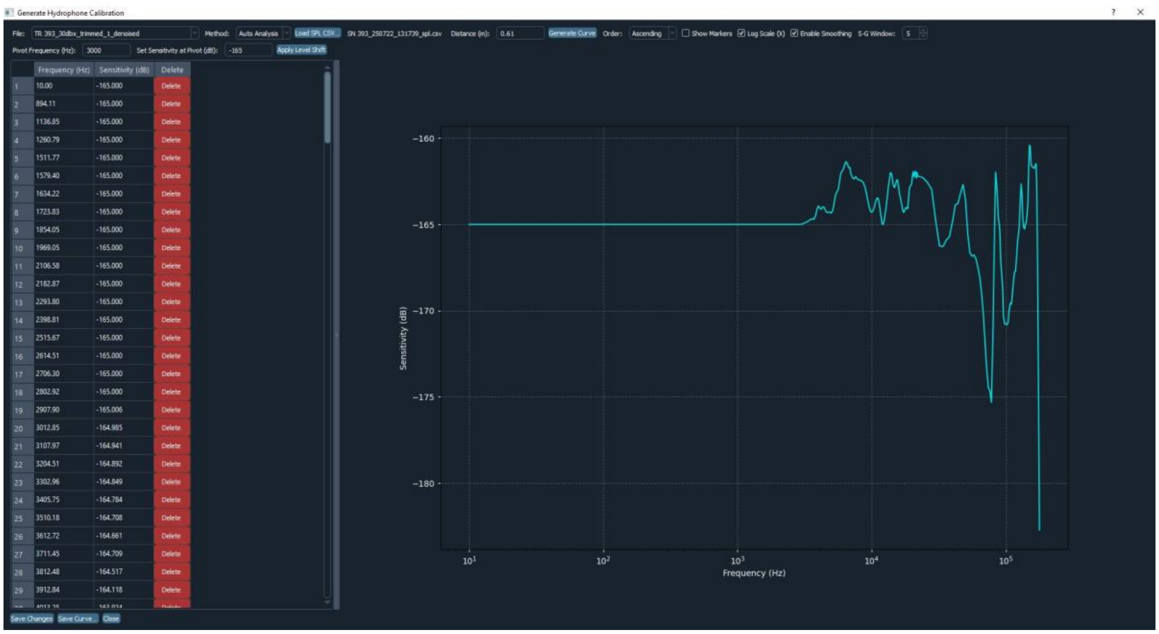

Figure 38: Hydrophone Calibration Popup

This tool builds and refines your hydrophone's sensitivity curve (dB re 1 V/µPa @ 1 m) by combining measured voltages and a reference SPL CSV. You can interactively reorder, smooth, level-shift, and export the curve.

Use Cases

Derive a Sensitivity Curve

– Generate dB sensitivity vs. frequency from your measured pulses (target_frequency & measured_voltage) against a known SPL reference file.

Merge & Compare Calibration Runs

– Load multiple methods or files in sequence, overlay curves, and see run-to-run variation.

Level-Shift to a Reference Point
– Insert a known low-frequency anchor (e.g. 10 Hz) or pivot at a lab-measured value to correct systematic offsets.

Smooth Out Measurement Ripples
– Apply a Savitzky–Golay filter to suppress ± dB oscillations from narrowband anomalies while preserving overall curve shape.

Export & Archive
– Save the finalized curve back to the DB, export as CSV for plotting or for input into SPL Transmit Analysis.

Effects on Your Calibration Curve

Feature	Effect on Sensitivity Curve
Order (Asc/Desc)	Ensures frequency points and sensitivity values stay matched—avoids “flipped” pairs that misalign points.
Log-Scale X-Axis	Visualization only—no data change, but reveals octave/octave-band features more clearly.
Enable Smoothing (S-G)	Reduces point-to-point dB “wiggles,” trading off some fine detail for a cleaner, monotonic trend.
Pivot Level-Shift	Adds a constant offset so that the curve passes through your chosen pivot (frequency, sensitivity) point.
Low-Frequency Insertion	Appends a 10 Hz point at your pivot sensitivity if missing—fills in extrapolated low-end response.
Distance Correction	If you enter a distance ≠ 1 m, adds/subtracts $20 \cdot \log_{10}(\text{distance})$ to all sensitivity values, converting between near-field and far-field references.

• Duty Cycle

Measures how much of the signal exceeds a threshold:

$$DC = \frac{T\ Active}{T\ Total} * 100\%$$

Useful for pulse trains and intermittent signals.

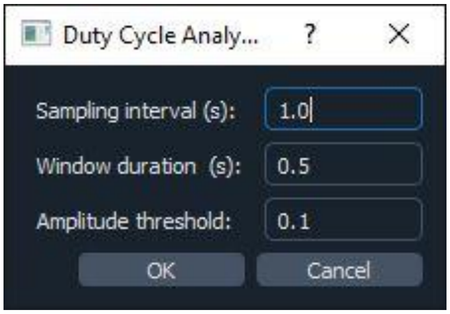

Figure 39: Duty Cycle Analysis Popup

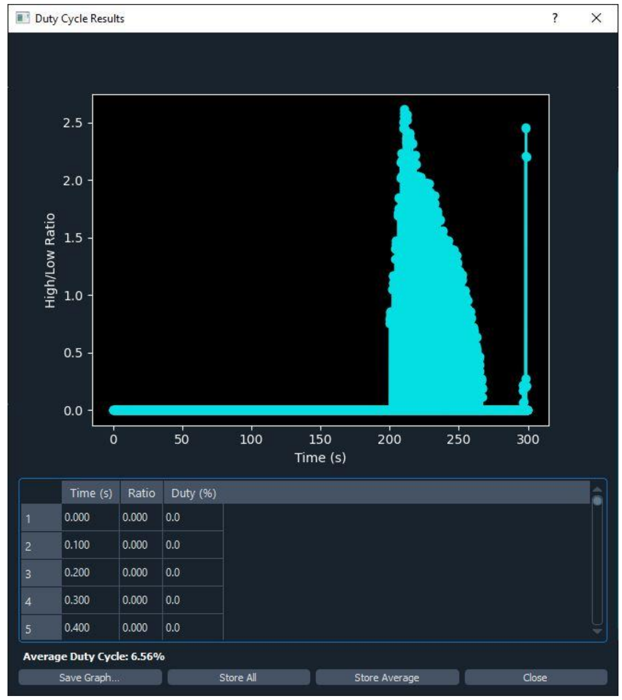

Figure 40: Duty Cycle Analysis Popup Results

The Duty Cycle tool lets you quantify what fraction of a recording lies above a user-defined amplitude threshold—essentially, how “active” your signal is vs. “quiet” time. You pick a threshold, run the analysis over your waveform or selected region, and it computes and plots the percent-time above that level.

Use Cases

Pulse Duty Characterization

– Measure the on-time vs. off-time of pulsed sonar or LFM/HFM signals to verify transmitter duty cycles.

Ambient Activity Monitoring

– Track how often ambient noise (e.g. snapping shrimp, vessel noise) exceeds a safety or detection threshold.

Equipment Stress Testing

– For continuous machinery noise, measure percent-on time to assess duty loads on hydrophones or amplifiers.

Adaptive Thresholding

– Use duty-cycle output to set dynamic detection thresholds for downstream pulse or click analyzers.

Effects on Measurement Values

Logged Metric	Effect of Duty Cycle Settings
Duty Cycle (%)	**Direct output**: percent of samples above threshold during analysis window.
Threshold Level	Raising threshold → **lower** duty% (fewer samples qualify); lowering → **higher** duty%.
Window Selection	If you analyze a subset (region), duty% applies only there; whole-file analysis gives overall.
Sampling Rate	Higher fs → more precise time quantization of on/off transitions → slightly different duty%.
Time Resolution	Without windowing, uses per-sample. If a sliding window is later applied, window length smooths short spikes.
Logged Misc Value	Stored in `misc` column as duty-cycle fraction or percent—make sure you note units in your metadata.

5. Detection & Classification Tools

• Active Sonar Detection

Performs matched filtering to detect sonar-like signals.

Matched Filter Principle:
\[
y(t) = x(t) * h(-t)
\]
Where h(t) is the time-reversed version of the known transmit signal.

This boosts SNR for signals that match the known pattern.

Results include:
- Time of detection
- Matched-filter output magnitude
- Estimated range (from time-of-flight if applicable)

The Active Sonar tool lets you configure a transmit–receive experiment: you specify the ping waveform, listen for echoes, and log range, amplitude, and arrival-time metrics for each echo return.

Use Cases

Depth Profiling & Bathymetry
– Send a single ping and measure the time-of-flight to the seafloor or object to compute depth or range.

Fish & Object Detection
– Emit short pulses and capture multiple echo returns from fish schools, obstacles, or wrecks—logging each return's strength and delay.

Impedance & Target Strength Calibration
– Use known reflectors (e.g. calibration sphere) to measure absolute echo amplitude vs. range for system calibration.

Scan-Mode Surveys
– Automate a sequence of pings at fixed intervals or along a survey track; collect echo series for successive spatial sampling.

• Cepstrum Analysis

Used to detect periodic echoes (e.g., reverberation, multipath).

Peaks in cepstrum correspond to time delays between echoes.

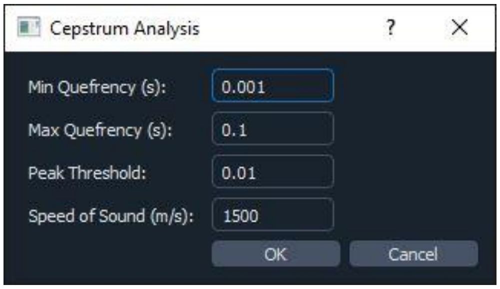

Figure 41: Cepstrum Analysis Popup

Figure 42: Cepstrum Analysis Popup Results

The Cepstrum tool computes the cepstrum (the inverse FFT of the log-magnitude spectrum) of your waveform, letting you isolate periodic structures—such as echoes, multipath delays, or the spectral envelope. You can apply **liftering** (cepstral filtering) to emphasize or remove these components.

Use Cases

Echo & Multipath Delay Estimation

Identify the time delay between direct arrivals and reflections by reading the **quefrency** peak corresponding to round-trip delays.

$$\tau_p = \frac{n_p}{f_s}$$

where f_s is the sampling rate. In an underwater echo context, $\tau_p$is the round-trip travel time of the reflection.

Spectral Envelope Extraction (Deconvolution)

Separate system/hydrophone response (slow-varying envelope) from fine spectral detail (fast-varying), useful for flattening response before SPL or FFT analysis.

Pitch / Fundamental Frequency Detection
In tonal signals (e.g. marine mammal calls), the quefrency domain reveals pitch periods more clearly than the spectrum.

Noise & Reverberation Suppression
By zeroing out low-quefrency components (liftering), you remove slowly-varying reverberation or hum, improving transient detection.

$$C_{liftered}[n] = \{\frac{0,}{c[n],} \frac{0 \leq n \leq n_1}{n > n_1}$$

• Event Clustering

Performs unsupervised clustering of detected events using k-means.

Steps:
1. Extract features (e.g., duration, peak freq, amplitude)
2. Apply PCA if necessary
3. Run k-means to group similar events

Useful for:
- Grouping marine mammal call types
- Categorizing noise bursts

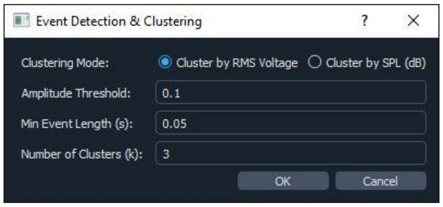

Figure 43: Event Detection & Clustering Popup

	Start(s)	End(s)	RMS Voltage (V)	DomFreq(Hz)	Cluster
1	0.336	0.442	0.1884	7769.53	2
2	0.442	0.546	0.2827	7740.23	2
3	0.571	0.629	0.1827	7675.78	2
4	0.687	0.739	0.1984	7558.59	2
5	0.739	0.861	0.2077	7494.14	2
6	0.861	0.921	0.1752	7476.56	2
7	1.662	1.771	0.0813	5835.94	1
8	1.888	1.958	0.1277	5894.53	1
9	1.987	2.053	0.1330	8578.12	2
10	2.152	2.226	0.1021	8742.19	2
11	2.301	2.357	0.0382	9046.88	2

Figure 44: EVent Detection & Clustering Popup Results

The Event Clustering tool groups detected acoustic events (peaks, pulses, clicks) into meaningful clusters based on time proximity, frequency similarity, or other attributes—helping you identify trains of clicks, burst calls, or episodic noise events.

Use Cases

Bioacoustic Click Trains

Group individual clicks into dolphin or porpoise trains by clustering on inter-click interval (ICI).

Burst-Pulse Sequences

Identify series of pulses (e.g. sonar chirps, machine knock sequences) that occur as part of a single transmission.

Ambient Event Episodes

Cluster short-duration noise bursts (e.g. seismic airguns, piling noise) into episodes for aggregate SPL or duty-cycle metrics.

Outlier & Noise Removal

Group sparse, isolated noise peaks into a “noise” cluster that can be filtered out of analysis.

Effects on Logged Metrics

Logged Metric	Effect of Clustering
Event Count	• Individual events are grouped—total "clusters" count decreases relative to raw peak count.
Cluster Duration	• Time span from first to last event in each cluster—useful for burst-length statistics.
Inter-Event Intervals	• Within-cluster ICI statistics (mean, std) can be computed instead of per-event timing alone.
Aggregate RMS/SPL	• You can compute a single RMS or SPL per cluster window rather than per event—smooths variability.
Duty Cycle	• Duty cycle can be computed on cluster windows (time "active" per cluster), giving a burst-level duty metric.
Frequency Stability	• Clustering by frequency allows you to track runs of events at similar frequency—detects shifts or drifts within a cluster.

6. Database Tools

The application uses a local SQLite database (`analysis_log.db`) to store all measurements, classification results, and SPL analysis data. Users can manage, review, and export results through dedicated database tools.

• Log Viewer & Filter

Displays all entries in the database with options to filter by:
- File name
- Tool name or category
- Date/time range
- Frequency range

Results are shown in a scrollable table view.

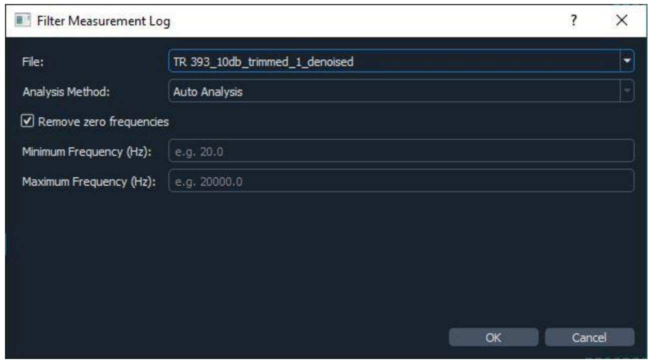

Figure 45: Measurement Log Viewer

• Export Logs to Excel

Allows export of selected log entries into an `.xlsx` file for offline use or reporting.

Exported data includes:
- File name
- Tool/method used
- Measurement values (e.g., RMS, SPL)
- Timestamps
- Frequency
- Comments or annotations

• Annotate Results

Users may label results with custom tags such as:
- "Whale Detected"
- "Bad Hydrophone"
- "Clipped"
- "Reference Tone"

Annotations are saved in the database for filtering or export.

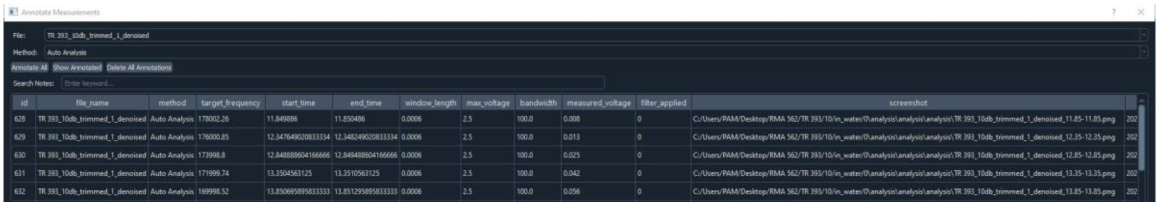

Figure 46: Measurement Annotation Viewer

• Delete Log Entries

Supports deletion of:
- Selected rows
- All results from a specific tool or project
- Frequency-specific deletions

Useful for re-running tools or managing batch processing outputs.

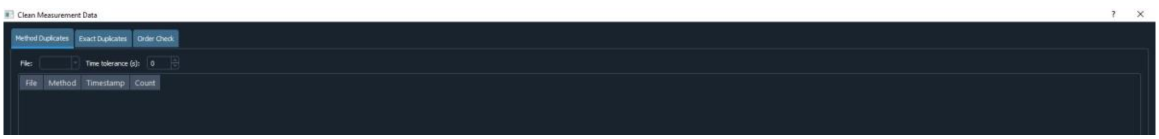

Figure 47: Clean Measurement Data Method Duplications

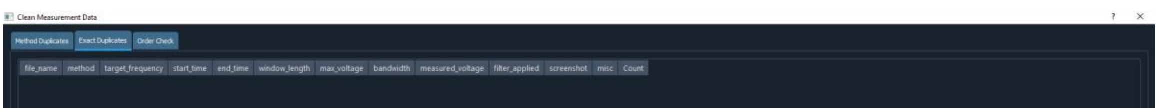

Figure 48: Clean Measurement Data Event Duplication

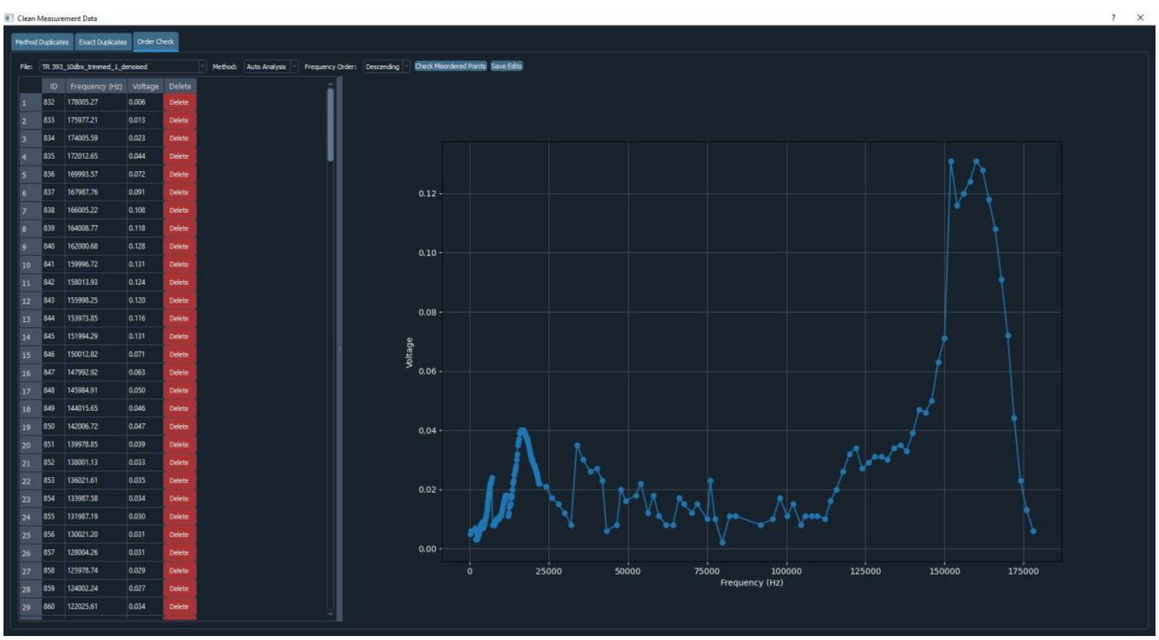

Figure 49: Clean Measurement Data Order Check

• Voltage Correction Tool

Adjusts voltage-based measurements (e.g., SPL, RMS) by applying a known correction factor (gain or attenuation).

Correction applied as:

$$V_{corrected} = V_{measured} * K$$

Where:
- K is the correction factor (e.g., due to probe scaling, user input error, or hardware calibration mismatch)

Figure 50: Voltage Correction Popup

• Calibration Curve Manager

Allows creation, editing, visualization, and export of hydrophone or transducer sensitivity curves.

Features:
- Add new curve (name, frequency, sensitivity pairs)
- Edit/delete frequency rows
- Overlay multiple curves
- Export to CSV
- Plot with dark/light mode

Sensitivity values (in dB re 1V/µPa) are applied during SPL calculations.

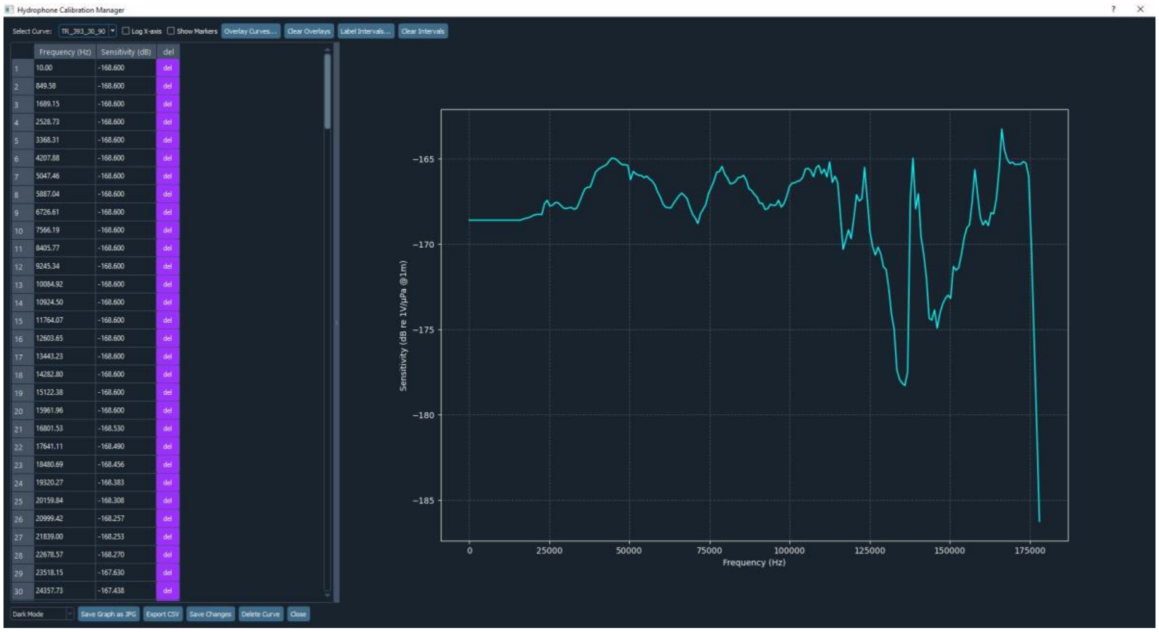

Figure 51: Hydrophone Calibration Curve Manager

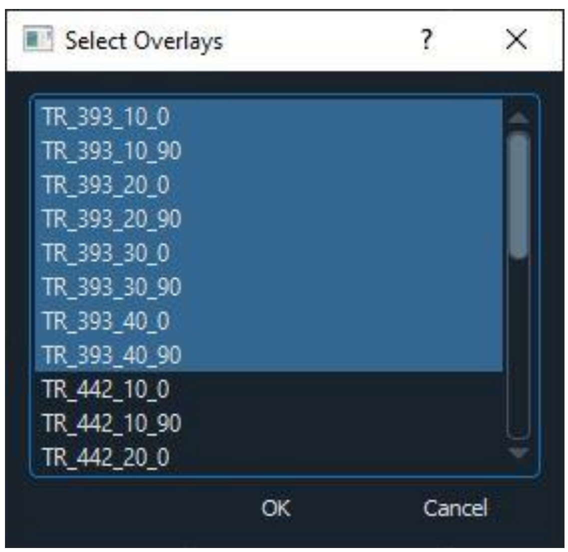

Figure 52: Hydrophone Calibration Curve Manager Overlay Popup

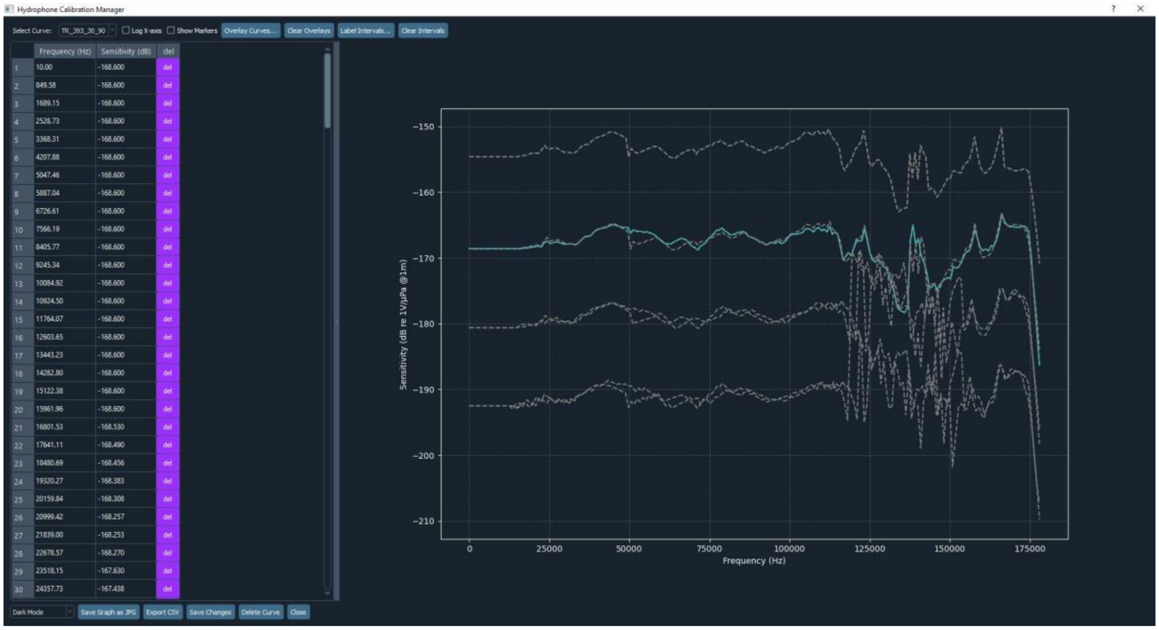

Figure 53: Calibration Curve Manager Overlay Popup

7. Projects Tab (File × Tool Matrix)

Organizes files under named "Projects". Each file × tool entry shows:
- ✅ View – results already exist
- ▶ Run – tool not yet executed

Supports all tools: measurement, classification, SPL, waveform metrics.

Use Case:
- Monitor progress on large batch analysis
- Track what has been completed across deployments

Figure 54: Projects Tab

The **Projects** tab gives you a one-glance dashboard of all your analysis and development efforts. Below is a more in-depth walkthrough of its layout, controls, and how to get the most out of each element.

8. SPL & Calibration Tab

This tab displays Sound Pressure Level (SPL) results derived from RMS voltage values and user-provided calibration curves.

Calibration Process:

1. Import a known SPL vs frequency reference.
2. Record the waveform from the system under test.
3. App extracts RMS voltage per tone.
4. Sensitivity is computed:

$$S(f) = SPL - 20 * LOG_{10} V_{rms}(f)$$

Supports:

- Per-tone calibration
- Overlaying computed and reference curves
- Storing curves for reuse

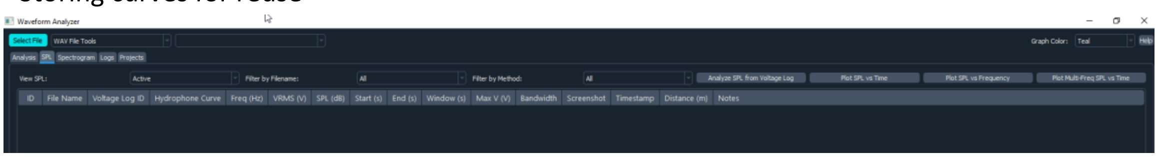

Figure 55: SPL Tab

9. Spectrogram Tab

Parameter Controls (Top Row)

Control	Purpose & Details
Threshold (Amp):	Absolute amplitude above which an event is detected. • *Default:* 1000 sample units. • Increase to ignore low-level noise; decrease to catch fainter events.
Pre-buffer (s):	Seconds to include **before** the first above-threshold sample. • Captures the leading edge of pulses or artifacts.
Post-buffer (s):	Seconds to include **after** the segment's minimum length. • Ensures you capture full echo tails or decays.
Min Length (s):	Minimum duration of each extracted segment. • Short detections are auto-extended by the post-buffer.
Spectrogram NFFT:	FFT frame size for spectrogram computation. • Options: 512–16384. • Larger values → better frequency resolution, poorer time resolution.
Colorbar:	Toggle on/off the decibel colorbar legend. • Off to maximize plot area; on to see amplitude scale.
Colormap:	Select Matplotlib colormap (e.g. inferno, cividis). • Influences how you perceive dynamic range.
Theme:	**Dark** (black BG, white axes) or **Light** (white BG, black axes). • Matches your overall UI preference.
Min (Hz):	Lower Y-axis (frequency) limit. • Helps zoom in on the band of interest (e.g. 1000–5000 Hz).
Max (Hz):	Upper Y-axis limit; blank defaults to Nyquist (fs/2). • Suppresses ultrasonic or out-of-band noise.
Generate Spectrograms:	Triggers detection and segmentation. • Scans the entire WAV for events above threshold. • Populates spec_segments and displays the first segment.

Spectrogram Display

- **MplCanvas**
  - A Matplotlib figure embedded in the tab, sized 8×6” at 100 dpi.
- **Axis Configuration**
  - X-axis: time in seconds (relative to the segment start).
  - Y-axis: frequency in Hz, limited by **Min/Max (Hz)**.
  - Tick marks, labels, and spines adapt to the selected **Theme**.

SpanSelector (Region Selection)

- **Function**: Click-and-drag horizontally on the spectrogram to define last_spec_region = (xmin, xmax).
- **Use**: Determines the precise time slice for **Save ML Clip**.
- **Visual Cue**: A semi-transparent overlay appears over the selected region.

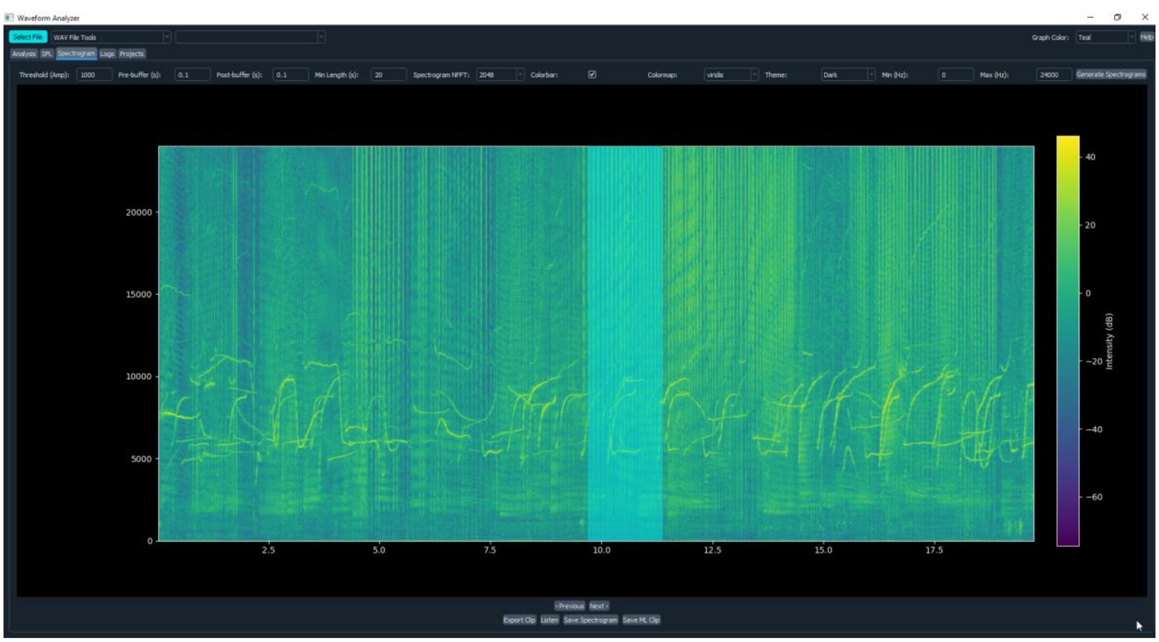

Figure 56: Spectrogram SpanSelector

Navigation Buttons

Button	Action
‹ Previous	Go to the previous detected segment. Disabled when at the first segment.
Next ›	Advance to the next segment. Disabled when at the last segment.

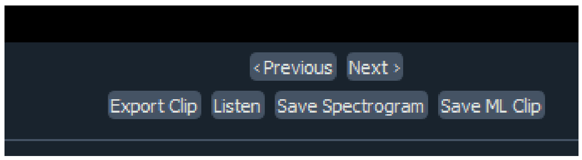

Figure 57: Spectrogram Navigation Buttons

Export, Listen & Save Buttons

Button	Enables When…	Action
Export Clip	After generation (any segment)	Saves the **entire** current segment (pre + event + post) as a WAV file via file dialog.
Listen	After generation	Plays back the current segment immediately through the system audio device.
Save Spectrogram	After generation	Exports the current spectrogram image as JPG/PNG (150 dpi), preserving theme & colormap.
Save ML Clip	After generation **and** region selected	Exports only the SpanSelector-defined slice from the current segment into ml/<label>.wav.

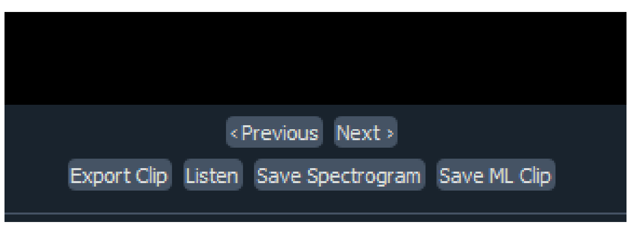

Figure 58: Spectrogram Export, Listen, & Save Buttons

Typical Workflow

1. **Set Parameters**
   – Choose your amplitude **Threshold**, **Min Length**, and **Buffers** for pulse capture.

2. **Generate**
   – Click **Generate Spectrograms**; the WAV is scanned, and segments are extracted.

3. **Navigate**
   – Use **‹ Previous** / **Next ›** to review each segment's spectrogram.

4. **Select Region (Optional)**
   – Drag on the spectrogram to highlight a sub-region for ML training data.

5. **Export or Listen**
   – Use **Export Clip** or **Listen** to verify the audio content.

6. **Save Outputs**
   – **Save Spectrogram** for publication-quality figures.
   – **Save ML Clip** for machine-learning datasets.

Tips & Best Practices

Threshold Tuning

Why: Helps users choose a data-driven threshold rather than guessing.

Where: In the "Parameter Controls" section, immediately under the threshold description.

Choosing NFFT

Why: Different NFFT sizes dramatically change your time vs. frequency clarity—critical for getting interpretable spectrograms.

Where: In the "Spectrogram NFFT" entry, as a bulleted note on recommended sizes.

Buffer Settings

Why: Proper pre- and post-buffers ensure you capture the full event without clipping the start or tail.

Where: Next to "Pre-buffer" and "Post-buffer" descriptions as minimum suggested values.

Frequency Limits

Why: Suppresses out-of-band noise (e.g. ultrasonic or digital artifacts), making your plots cleaner.

Where: Under "Min (Hz)"/"Max (Hz)" controls, as a reminder to set "Max" just below your hydrophone's upper spec.

Colorbar Usage

Why: Maximizes visual real-estate in presentations while still providing quantitative amplitude info when needed.

Where: In the "Colorbar" toggle description—note when to disable vs. enable.

Theme Selection

Why: Dark vs. light dramatically affects legibility in different contexts (screen vs. print).

Where: In the "Theme" control entry, with guidance on use cases.

Colormap

Viridis, Plasma, Inferno, Magma, Cividis

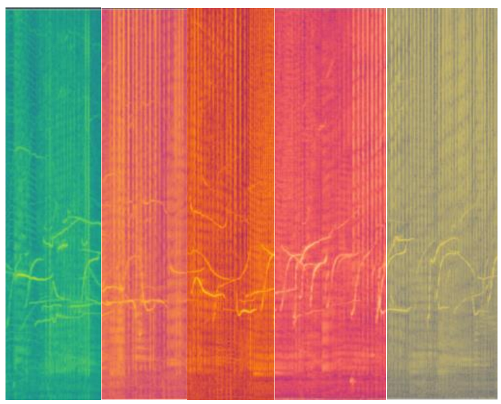

Parameter Controls

Control	Description
Threshold (Amp)	Amplitude threshold (sample units). Segments where absolute sample value ≥ threshold are extracted.
Pre-buffer (s)	Seconds of audio before the first above-threshold sample, capturing waveform onset.
Post-buffer (s)	Seconds after the minimum segment length, ensuring you capture the full pulse or echo tail.
Min Length (s)	Minimum duration of each segment. Shorter events are extended by the

Control	Description
	post-buffer.
Spectrogram NFFT	FFT frame size (samples). Large NFFT → better frequency resolution, poorer time resolution; small NFFT vice versa.
Colorbar	Toggle the intensity legend on or off.
Colormap	Matplotlib colormap for the spectrogram (e.g. `viridis`, `inferno`, `plasma`).
Theme	**Dark** (black background, white axes/text) or **Light** (white background, black axes/text).
Min (Hz)	Lower frequency limit on the Y-axis (Hz).
Max (Hz)	Upper frequency limit; leave blank to default to Nyquist (fs/2).
Generate Spectrograms	Scan the full audio for events above threshold, extract segments, and display the first spectrogram.

---

10. Licensing System

Waveform Analyzer uses a hardware-locked license mechanism. The license mechanism ensures that users are not copying the application and installing it on unauthorized machines. If you require more than one license, please contact sales@blackfishacoustics.com!

Steps:

1. App requests user information:

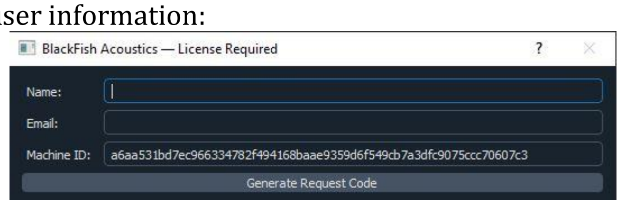

Figure 59: License Request Popup

2. App generates `license.req`
3. User sends file to sales@blackfishacoustics.com
4. Receives `license.dat` file via email. Save license.dat to the application directory.
5. App verifies file signature and hardware match.

If no valid license is found:
- User prompted again for details to generate the license.req file.
- App exits if verification fails

Common license issues:

1. Hardware ID changes due to system changes.
    a. Changing hardware configuration, like installing new hard drives or graphics cards, can cause the hardware ID of your machine to change. If this occurs, contact sales@blackfishacoustics.com to receive a new license file.
2. License file is named wrong!
    a. If you change the name of the license.dat file, the application will not see the file and will not load the license.
3. License file has expired
    a. If the license file was purchased with a certain amount of time for activation, when the license date arrives the application will not function. In this case, you must contact support to purchase a new license.

11. Help & Troubleshooting

Help access: Click Help in the toolbar.

Common Issues:
- Blank spectrogram: Verify WAV data and length
- SPL 0 dB: Calibration curve may not span entire frequency bandwidth
- Tool doesn't launch: Project or method mismatch
- License error: `license.dat` may be missing or invalid

For any user assistance, please email sales@blackfishacoustics.com