LYZ3R ROOM
Help and FAQ

Record, inspect, document and save sound on your phone.

Use this guide to find what a control does and what it changes. The same topics are searchable offline inside the app. Original recordings, measurement references and file associations matter: keep the take package together.

Contents

01 Start here

02 VOICE REC speech recordings

03 SLM sound level meter

04 3CH0 impulse and echo analysis

05 REC0R recording and live views

06 Live contrast and plot icons

07 LYZ3R source and analysis

08 LYZ3R zoom and display settings

09 Photos and video with measurements

10 Replay crop and save packages

11 GEN3R8 sources and output

12 GEN3R8 Sweeps and Draw

13 SL33P scenes and timer

14 Take information GPS and environment

15 Frequently asked questions

16 Illustrated quick reference

Guide revision 2026-10-30 FFT controls in plot settings and Dynamic for app 0.4.44. The experimental level meter is not certified. Features described here accompany this consolidated build; older installations can differ.

01 Start here

Choose the job first, then use the controls in that module. Help works offline; only the optional nearby-weather lookup needs the internet.

Make your first recording

Choose REC0R from the ROOM module selector if another startup module is selected. You can also choose it from the LYZ3R ROOM module selector. Leave the camera closed for an audio-only take. Tap the red Record button, make your sound, then tap Stop. Replay checks the take; Save keeps a recording package. Analyze opens its sound in LYZ3R. Save complete dataset retains full original WAV, original camera media when available, notes and JSON for later processing. Export to share writes only the selected finished copies, without a raw-data package. Export does not archive originals; save the dataset before replacing a take if you may need it later.

ROOM module menu and Arrange apps

Tap the full LYZ3R ROOM logo and current module name to open the module menu. Arrange apps lets you change menu order and choose exactly one Open on launch module. The saved startup choice and menu order are independent. REC0R is the fallback until you choose another module. Opening the app never starts recording or playback automatically. Switching modules stops playback and capture; audio-only background recording follows the screen-off preference. Android Back closes the current layer before opening the module menu.

Back arrow and Android Back

Back closes the current help, settings, save or replay layer first. It can restore an enlarged plot or leave camera view. At the main workspace it opens the module selector. Closing Save cancels that save operation; it does not mean the take has been saved.

Moon, sun and information

The moon/sun button changes the colour theme. The information icon opens the current module notes. Help & FAQ opens this searchable guide. Secondary information can also be reached by an inward swipe from a screen edge in SLM, REC0R and LYZ3R. Use the visible buttons if a phone gesture takes priority. Info and Help show the installed release version. Quote this version when reporting a problem. A test candidate is identified explicitly; the version does not mean the app is certified.

Choose your phone folder once

Open the module menu, then Settings and Choose recording folder. For the consolidated phone files, select Documents / MEQ App and allow Android folder access. Keep one folder per take inside it. Related audio, video and information belong together. Changing or moving the folder requires selecting its new location again.

Which app does this guide describe

This guide matches the consolidated Android build including normal REC0R's two-window layout, live Dynamic contrast and Waterfall, plus LYZ3R's analysis views. Older installed versions can differ. It documents implemented behaviour, not certification or a guarantee for every Android device. Module purchasing is not active.

Search without the exact button name

Help search works locally across modules and accepts partial words, common typing errors and a curated set of German/English synonyms. Try speichern for saving, Kamera for camera, or replay for playback. It ranks matching titles first. It is not a cloud assistant or audio/image-content search. Recording import search separately filters titles, dates, locations and notes in the chosen folder. Empty workspaces hide actions that need existing data. Recording and mode/settings controls remain discoverable; relevant actions return when a take is available.

Draft and send feedback to MEQ

Open Info, then Feedback to MEQ. Write a subject and feedback; the draft stays on this device. You can include the app version and current module, copy the feedback, save a TXT draft, open an email draft to info@meqtechnologies.com, or open the MEQ contact page. To use the contact form, first copy the feedback, open the page, paste into the message field and send there. Opening a page or email draft does not send your message. No recordings are attached automatically. If no email app is available, use Copy feedback and the contact page.

Set FFT block size and Hann / Hanning window

In Spectrum, Spectrogram or Waterfall, open Plot settings or Dynamic to set FFT block size (1,024, 2,048, 4,096, 8,192 or 16,384 samples) and window function (Hann / Hanning, Hamming, Blackman or Rectangular). The choices stay synchronized across these controls and tools. Larger blocks improve frequency bin spacing but span more time. The settings show bin spacing and block duration for the current sample rate. Hann reduces leakage around tones; Rectangular is most suitable for coherent periodic blocks. Settings change the actual spectral calculation, including camera overlays, without changing the recording or SLM level calculations. Live spectral history restarts when its FFT settings change so incompatible bins are not mixed. Open recordings in LYZ3R to recalculate their full selected interval. Channel waveforms do not use an FFT.

02 VOICE REC speech recordings

A simple speech recorder with local transcription and optional external translation. Keep original audio, editable transcript and translation as separate copies.

Record, Stop and retained recordings

Choose VOICE REC immediately after 3CH0 in the module menu. Record starts audio capture; the same button becomes Stop. Audio is retained before transcription. Audio-only capture continues in retained four minute WAV parts until stopped or low storage. Open continuous recording parts reopens completed parts; the final partial part opens after Stop. Open audio imports a PCM WAV up to 64 MiB. Manage retained recordings deletes chosen Voice copies and drafts; native recording parts and exported files remain separate. Native microphone permission is requested only when capture starts. Voice recordings and text drafts are retained on this device; the Retained recordings list reopens prior audio. Save copies into your chosen MEQ folder for durable access outside the app. Screen-off behaviour follows the normal recording preference; camera recording is not part of VOICE REC. While recording, the moving waveform and independent input meters show the captured signal. The red overload indicator holds for three seconds; thin red waveform rail regions locate digital clipping. The Stop button stays reachable while you scroll. Import latest REC0R take brings the retained native recorder audio into a new voice draft without changing previous retained voice copies. Clear returns to an empty workspace, including after reopening Voice Recorder. Prior audio and text remain in Retained recordings and saved files are unchanged. Clear is unavailable during recording or transcription.

Choose speech sample rate and input

The Setup button with sliders icon opens a microphone selector and 16, 24 or 48 kHz recording rate. VOICE REC defaults to 16 kHz, the lowest offered speech setting. Higher rates retain more audio detail and use more space. REC0R keeps its independent high-quality preference. Settings apply to the next take; the actual native rate and mono/stereo state appear while recording. Unsupported input/rate combinations fail visibly. USB stereo still preserves both original channels; choose Left or Right for transcription.

Transcribe a minute or longer

Stop the recording, open Transcript, choose the spoken language and tap Transcribe. A bundled quantized multilingual Whisper Tiny engine processes a 16 kHz speech copy locally in a worker with overlapping long-audio chunks. Audio is not uploaded and no account or network is required for transcription. The UI reports the actual preparation/processing stage; there is no invented progress percentage. Cancel leaves the recording and existing text intact. Long inference can take time, especially on a phone. Read and edit the draft against replay; names, numbers, soft speech and background noise can cause recognition errors. There is no automatic speaker identification, summary or translation.

Mark, crop and preview voice audio

Voice uses the Channel gestures: drag directly on the waveform to mark, refine the outward edge triangles, or move the complete interval using the rounded three-line centre grip. Two fingers zoom and pan without changing the mark; +, minus and Fit view also change only the visible window. Start and End accept exact times. Play marked auditions only the marked sound. Keep marked immediately keeps that interval; Remove marked immediately removes it and joins the remaining sound. The waveform, player and duration update together. Undo reverses the last edit; Restore original returns to untouched audio. Select all changes the mark, not the view or audio. Save and Transcribe use the displayed working audio. Original audio, edits, selection and zoom are retained when reopening. Actual audio edits flag an older transcript as out of date; moving marks or zooming does not. Saved originals are never overwritten.

Save audio and complete text files

Save transcript writes the entire editable transcript as UTF-8 TXT, avoiding clipboard length limits. Save audio + transcript writes both through the existing chosen-folder workflow. A cancelled or failed second save does not erase the first saved file. Save translation writes a separate language-labelled TXT. The original WAV and original-language transcript stay available independently from translation. No save button sends your content to MEQ.

Translate with Google Translate

Translate opens the installed Google Translate app with your transcript, or the selected text if you highlighted a section. The visible helper explains: choose a language there, copy the result, return here and tap Paste translation. The translation has its own editable field. Clipboard content is read only when you tap Paste translation, never automatically on return. Existing translation edits require confirmation before replacement. Choose an optional language label and Save translation (.txt). If Google Translate is unavailable, use Copy text or Share text. Long content is never silently truncated: select a shorter section or save the full text file. Google Translate controls its own online/offline operation and translation; this is external-app handoff, not a paid translation API or automatic return.

03 SLM sound level meter

Use SLM to watch changing sound levels. You can start immediately; a reference reading is optional.

Record, Pause, Resume, Stop and Clear

Record captures microphone audio and sound-level measurements together. Pause stops adding samples and duration; Resume continues the same measurement. Stop releases the microphone and retains the result. Record after Stop creates a new session while retaining previous sessions under Open. Clear empties the current view and never deletes previously retained or exported sessions. It is disabled while recording. Weighting, response and input settings stay fixed for a session; clear the current view before changing them.

Current level, Average, Maximum and Time

The large number is the current time-weighted level. Average is the energy average, Leq, over the measurement. Maximum is the highest time-weighted level since the session began; its marker also appears on the dial/history. Time is the session duration. These quantities answer different questions and need not have the same value. Average, Max and Time share a compact row. Units and weighting follow the main reading and the A/C/Z and Fast/Slow indicator. Average, Max and Time share a compact row; the main readout shows the units and the header shows weighting and response.

Weight and Fast or Slow

Weight cycles A, C and Z. A reduces low-frequency contributions; C is flatter; Z is unweighted. Fast uses a 0.125-second power response and Slow uses 1 second. Changing weighting or response starts a fresh measurement. The ordinary spectrogram remains unweighted even when the main meter uses A or C. The large level and history labels show dBA, dBC or dBZ. Estimated units retain est. Digital input remains dBFS. Fast (125 ms) and Slow (1 s) describe the time response and maximum; the energy average is independent of this response. The narrowband spectrogram remains explicitly Z weighted.

History and Spectrogram

History shows level against time. Spectrogram shows changing frequency content using a 2048-sample Hann analysis and a logarithmic frequency axis, approximately 50 Hz to 20 kHz within the source limits. The live camera also offers independent Channel, Spectrum and Spectrogram overlays. Select more than one when useful. The spectrogram has a vertical colour scale on its right; higher levels are at the top. Dynamic is beside the plot tabs. Small screens scroll rather than reducing the plot below usable height.

Estimated, Reference and Setup

Estimated means a generic microphone sensitivity is being used; it is not a calibrated acoustic reading. Setup contains frequency/time settings, input choices and Adjust reading. Optional comparison with a reference meter or a compatible calibrator applies a scalar adjustment. Recheck it after changing microphone, input or gain. This experimental meter has unverified phone processing and acoustic accuracy; it is not a certified survey or exposure instrument.

Enlarge and PNG

The outward-arrow button enlarges the plot; Back restores it. PNG saves the displayed graph with units and reference/estimation information, not the original audio. Use a camera recording and Save package when you need original audio and video for later analysis. PNG opens an exact image preview first. Save PNG writes that preview; Back returns without saving.

Save and reopen sound with complete measurement history

Save prepares a .slm session and a separate PDF report in the selected package folder. A .slm file is a lossless ZIP container holding original PCM WAV, complete measurement history, settings, reference information, CSV and report PDF. The live history may show only the latest minute, but retained and exported history covers the complete session. Open lists locally retained sessions, including interrupted ones, and can import a .slm file or supported WAV. A restored measurement opens its saved history and audio together. Touch the history to seek that part of the sound; playback shows its position in the history. Save never clears the retained result. Cancelled or failed exports retain the source for retry. Plain WAV imports are explicitly estimated; a WAV without its associated reference cannot establish the original acoustic SPL.

Initial sound investigation report

Save includes a short PDF with date/time, duration, average Leq, time-weighted maximum, full-session history graph, microphone identity, weighting/response, reference status, notes, location and interruptions. It includes an audio checksum for integrity checks, not proof of the event or timestamp. It is an initial investigation report, not a certificate of regulatory compliance. Original sound, full observations and complete event list remain in the session file. Minimum, true acoustic C-weighted peak, dose and occupational exposure assessment are not claimed.

Save microphone identity and reference notes

Setup offers microphone profiles for a model, serial/identity and the current reference record. Select the connected input and perform an appropriate reference check for the actual phone, gain and microphone. Loading a profile does not apply an old calibration automatically. A profile or external microphone alone does not make this system a certified sound level meter. Previously saved sessions keep their original setup information when reopened.

Storage, interruptions and capture limits

On Android, SLM records microphone sound through a foreground service with an ongoing Stop notification. Screen locking and switching to another app do not end the recording. Sound is checkpointed to a native WAV about once per second. Measurements process every captured sample, catching up after returning if the display was suspended. Pause excludes that interval. Stop finishes the audio and measurements before Save becomes available. Closing the app or a microphone interruption stops capture. Open offers native recovery WAVs if Android closed the app; these recalculate estimated levels, while saved measurements retain their original reference. Delete recovery sound removes that separate WAV. Browser capture still requires keeping the page visible. Camera video and reference adjustment require the app open.

04 3CH0 impulse and echo analysis

Inspect a recorded clap or impact using the same loaded source as LYZ3R. 3CH0 does not activate the microphone or duplicate the audio buffer.

How do I measure a reflection?

Choose 3CH0 immediately after SLM in the default module menu. Import a recording or use the latest stopped REC0R take. LYZ3R and 3CH0 share the source and time selection. Select an isolated clap with the teal time handles. Find arrival suggests the first sample above one quarter of the selected peak; check it yourself because the source and noise can mislead it. Drag the teal D and orange R lines or edit Direct arrival and Reflection in seconds. The orange R starts 50 ms after the suggested arrival as a handle to place, not an automatically detected reflection. The readout says Marker delay until you place R. Reflection delay is then their difference in milliseconds, not absolute travel time or distance.

How do I zoom closely enough?

Focus clap selects a short interval around the direct arrival and reflection. Plus and minus zoom around the direct arrival; All restores the full take. Pinch horizontally to zoom, then drag through time away from the marker lines. Drag near a line to move that marker. Markers stay at their source times when zooming and panning. If the direct arrival is outside the view, the decay curve is hidden; Focus clap or All brings it back. Find arrival explicitly chooses a new event in the selected interval. Expand fills the workspace; Back restores it. Replay shows a moving position line while playing. View contains replay speed and video replay; the paired video remains available through the shared analysis tools.

What does the decay curve mean?

The upper plot shows a peak envelope of the unweighted channel after removing its selected-window DC offset for analysis. The lower plot is its backward-integrated energy normalized to 0 dB at the direct arrival. Values are relative dB, not acoustic SPL. A dashed tail is within 10 dB of the estimated tail noise; it should not be treated as a reliable decay fit. Clipping and a low crest factor produce warnings. This first version analyses recorded transients: it does not remove the clap, microphone or source response, perform sweep deconvolution, or calculate RT60.

Does 3CH0 need extra audio memory or change my recording?

The same LYZ3R PCM buffer is reused. At most 2048 plotting buckets hold envelope and energy data, with roughly 40 KiB of additional numeric arrays; canvas memory is separate. PCM and replay remain unchanged. PNG saves the plotted envelope and decay with marker metadata. Save package retains the original source and the current 3CH0 marker settings inside analysis metadata; marker settings are recorded for reference and are not automatically restored by import in this first version.

05 RECOR recording and live views

Record the phone microphone, monitor it without a camera, then listen, analyze or save the take.

Record and Stop

Record starts the input shown by the Input button. Auto chooses a connected USB receiver in stereo, otherwise the built-in phone microphone in mono. Choose a specific input to prevent fallback to another device. USB uses one two-channel stream, with an original stereo WAV and synchronized mono Left and Right WAV files. The verified route must match the chosen device. Stop retains the take. Choose the next sample rate in Recorder settings; 48 kHz is the default. Audio-only capture uses retained four minute PCM16 WAV parts so earlier sound remains available during a long session. Voice Recorder offers Open continuous recording parts to open or delete completed parts. The current partial part is available after Stop. Low device storage stops capture safely. Save before another take replaces the current one.

Clear

Clear is at the far left of the bottom actions, beside Save. It asks before clearing the current workspace. Keep take returns without changing anything. Saved packages are not deleted. Save a take you need before clearing or recording again.

Channel, Spectrum, Spectrogram and Waterfall

Channel stays above exactly one selected Spectrum, Spectrogram or Waterfall window. Log / Linear changes frequency spacing. Expand fills the available workspace with the current plots; Restore returns to the normal layout. Video retains its independent Channel/Spectrum/Spectrogram overlays and no Waterfall. Each plot has a small downward triangle at its top right for Channel view. Choose Channels together or Channels separately. Channel and Spectrum overlay orange Channel 1 / Left and teal Channel 2 / Right when together. Separate stacks the channels with the same scales. Together Spectrogram and Waterfall show the mean energy of both channels, labelled L + R; this does not mix or change the source samples. Mono has no separate right microphone.

What does the 30-second window mean

The upper Channel strip and lower historical plot retain the latest 30 seconds as recording continues. New Waterfall slices enter at the back and move toward the front in its own window, with frequency across the front, time into depth, and colour/height indicating level. This display window does not cut the saved WAV. Spectrum is a current frequency snapshot. Expand the plot for more space or use LYZ3R to inspect the complete stopped recording.

Replay, Analyze and Save

Replay opens inside the recorder plot, without a separate replay window. Play, Stop and the seek slider stay in this panel. Open Replay options for the From/To range, speed, plots and crop/export controls. Close replay with the cross to return to the recorder views. Starting another recorder action closes replay; switching modules stops playback. Analyze carries the latest stopped take and its associated camera media to LYZ3R, even before saving. Save opens package choices. A manually imported LYZ3R source is kept when you merely switch modules; Analyze is the explicit action for loading the latest recorder take. PNG exports the currently selected lower plot; it does not save the WAV.

What changes the waveform size

The channel scale adapts to the signal. A larger drawing does not mean the WAV was amplified. The white camera waveform and spectral colours are display overlays; display contrast and scale never alter the stored source audio. Default measurement levels are estimated; digital dBFS is available in Settings. The two-channel waveform uses spaced vertical bars so the two envelopes remain distinguishable. Orange is Channel 1 / Left; teal is Channel 2 / Right. Both use the same automatic amplitude scale.

Live View before recording

Live View checks the selected microphone input without creating or replacing a take. With USB stereo, both moving waveforms and independent LED meters appear. Stop Live returns to the retained take. Phone mono is clearly identified and never duplicated as a second microphone. Headphone monitoring requires a connected headphone output; stereo monitoring listens to Left while both original channels remain recorded separately. Live View stops when the app leaves the foreground. Spectrum, Spectrogram and Waterfall all follow the same Live View microphone stream; choosing a display does not open a second microphone or save a take.

Recorder settings and sample rate

The small settings control in REC0R opens the same Recorder section as suite Settings. Choose 48, 24 or 12 kHz for the next audio-only take. 48 kHz is the default; Android may use 44.1 kHz if needed. Explicit 24/12 kHz choices require device support and otherwise fail without replacing the previous take. Lower rates limit frequencies below 12/6 kHz respectively and reduce file size. Continuous audio-only recording retains four minute parts; low device storage stops capture safely. Video uses its own capture rate; no saved file is resampled or modified by this setting. The Input button in the recorder readout selects Auto, Phone mono, or an available USB receiver stereo input. Input cannot change during recording, Live View or camera capture. Source selection opens the microphone choices; its title identifies the detected input. Sample rate is selected using the large 12, 24 or 48 kHz buttons in Settings. These apply to the next recording. The Channel settings include Show level display, enabled initially and remembered.

Can I record with the screen off?

On Android, audio-only REC0R takes continue with the screen off by default. The small hint below Record shows whether screen-off recording continues or stops, matching Recorder settings; camera mode instead reminds you to keep the app open. Start recording while MEQ is visible. Notification permission is optional: declining or disabling it does not prevent screen-off recording. Recorder settings has Continue audio recording with screen off (on by default), Keep screen awake while recording (off by default), and Lock-screen recording controls. The optional controls show Stop only while recording, silently. If controls are hidden, return to REC0R to stop; Android may still show the recording service in its running-apps list. Stop finalizes the current WAV; Save creates a complete take package. Switching MEQ modules or closing its task stops the take. The five-minute limit still applies. Video is foreground-only and keeps the screen awake during capture. SLM and live preview still stop in the background. Recordings never restart automatically after process death.

Channel LED bars and five-second overload check

Each channel has its own digital sample-peak LED bar, dBFS readout and red OVERLOAD indicator. The red segment activates at an absolute PCM16 sample of 32760 or greater (about -0.002 dBFS), and stays lit for three seconds after the most recent qualifying sample on that channel. A new event restarts only that channelâ€™s hold. The bar decays when input stops; the red hold remains visible for its remaining time. These are digital clipping warnings, not calibrated analog overload detection. A microphone or receiver can distort or limit before the PCM reaches full scale. In the Channel plot, thin green rails mark the upper and lower boundaries. Red rail regions locate original positive or negative digital clipping in time; display gain and weighting do not create these marks. In separate mode each channel has its own rails. Clipping is highlighted in bright red with a subtle glow and does not flash. Level LEDs progress from green to yellow to red. The last red LED holds for three seconds after source clipping. Bright red, glowing Channel rail segments identify clipped source regions; those regions replace the green rail.

06 Live contrast and plot icons

Use these controls to make quiet detail readable without changing what you record.

Dynamic and the 20, 40 or 60 dB choices

Dynamic appears for live Spectrogram, and for RECOR Waterfall. Off keeps a fixed colour scale. A selected span follows the recent spectral peak: 60 dB shows a wide level range, 40 dB gives more contrast, and 20 dB concentrates on stronger detail. The colours change; microphone gain, measured levels and WAV samples do not. The chosen span is remembered across live views. RECOR Waterfall retains each captured slice and moves it from back to front. Its height axis stays fixed while Dynamic changes colour contrast, so past peaks retain their height. Channel uses its separate automatic magnitude scale; Spectrum keeps its own level scale. In normal RECOR, the Log / Linear button changes frequency spacing without changing recorded sound or clearing history.

When should I turn Dynamic off

Use a fixed range when you need colours to mean the same numerical level across captures. In LYZ3R you can type fixed top and bottom limits for comparison. A dynamic plot can make two different sound levels look similarly bright, so read its scale as well as its colours.

Camera, outward arrow, unzoom and PNG

In live SLM/RECOR, Camera activates the camera preview. In LYZ3R the film-camera icon on a time plot controls video frames or attachment; it does not record a new video. Outward arrows enlarge a plot. The unzoom button restores the full view. PNG exports the visible plot image and its metadata.

Do I need the camera to see frequency plots

No. Normal RECOR always shows a Channel strip plus your selected Spectrum, Spectrogram or Waterfall window without camera access. SLM also has a normal spectrogram. Open Camera only when you want a visual record of the scene.

What do the spectrogram colours mean

They map spectral level to the labelled colour scale. The bright end represents higher levels and the dark end lower levels. Read the displayed top/bottom values and units; the image is not a photograph of sound pressure. Rainbow and Ember can be selected for analysis in View.

07 LYZ3R source and analysis

Inspect a saved or current take without requesting a new microphone recording.

Import and source information

Import opens the recording library when a phone package folder is configured, with readable take information and available previews. You can also choose other audio/video files. The source name and selected interval are shown above the plots; File information expands the available technical metadata. Keep the WAV, video and take JSON together for the best re-import result. Try demo starts a separate simulated source and clears the previous recordingâ€™s notes, location and reference information. File information opens beside the take name and expands to use the available width.

Choose any analysis views

Channel shows the waveform over time. Spectrum summarizes frequency-bin power across the selected interval. Spectrogram shows frequency against time, with level encoded as colour. Waterfall uses the same coloured 3D surface as REC0R. Later slices are at the back and earlier slices at the front; its time axis gives actual FFT frame-start times. Buttons activate views independently. Several active plots intentionally allow scrolling. One active plot fits the available screen with its main controls. Multiple active plots may scroll; no fixed Channel height is imposed. Expand still opens the larger view. Channel and Spectrogram share aligned time axes and matching linked markers, so the same horizontal position refers to the same instant in both plots. The top-right settings cog on Channel, Spectrum, Spectrogram and Waterfall selects Channels together or Channels separately independently for that plot. Orange is Channel 1 / Left; teal is Channel 2 / Right. Together spectra and waveform bars overlay the two sources; separate stacks them. Together Spectrogram and Waterfall use mean channel energy, explicitly labelled L + R. This is a derived display only and avoids phase cancellation. Mono disables the two-channel view choice. Plot settings use a cog. Show axis values controls transparent axis labels independently per plot and is remembered. PNG remains in the plot toolbar. Channel is a single linear time-history waveform, independent of spectral display limits. Mono is shown once; real stereo retains its original channels. The former duplicate overview and tiny disclosure triangle have been removed.

Channel and Weight

Channel displays the original signed waveform with the selected analysis weighting and level reference. Its centre line is zero amplitude, not a sound level. A stronger centre line and foreground grid remain visible over dense signals. Dynamic opens the same dB-span editor used by Spectrogram and Waterfall; upper/lower level endpoints can instead be fixed. Original audio is unchanged. The channel selector reads 1 / Left and 2 / Right for stereo, or 1 / Mono for a single channel. It continues to choose the channel used for cursor measurements, echo analysis and numerical exports. The plotâ€™s together/separate setting changes only which channels are drawn.

Replay without another window

Replay plays audio inside the analysis workspace. The playback cursor is shown during replay and hidden while paused, leaving the Channel interval grips clear. Dragging either the Channel or Spectrogram marker moves both to the same time. The bottom status bar shows the start, end, duration and energy-mean RMS level of the complete interval between both markers, independently for Left and Right, with the current weighting and level reference. The selected-interval readout remains visible while a recording is loaded. This readout does not change the teal interval selectors or the original audio. Replay stays available in the top toolbar while scrolling. View offers pitch-preserving speed changes. Video replay opens only when requested. Waterfall replay shows a translucent source-time plane and the selected spectral slice; its timestamp follows the same playback position.

Video frames beneath a time plot

Use the film-camera icon on Channel or Spectrogram, or enable Video frames beneath Channel in View. Frames provide time references; touch to enlarge and tap to open video near that moment. If there is no associated video, attach the matching source. A WAV alone cannot recreate a video that was never saved.

Waterfall rotation and playback time

Drag the Waterfall surface left or right to rotate up to 75 degrees each side of its default view. Drag up or down to tilt up to 15 degrees. Reset view restores the default orientation. On a keyboard, focus the plot and use the arrow keys; Home resets it. The right-edge grip still scrolls between stacked plots. During replay a translucent plane follows the original recording time and highlights the nearest analysed FFT slice at that position. It stays visible when paused and follows seeking from the shared Channel or Spectrogram marker. The highlighted slice has the existing FFT time resolution; it is not a new measurement. Later frames are at the back in the default view. Time, frequency, level limits, weighting and samples are unchanged by rotation. Normal REC0R Waterfall uses the same drag limits and Reset view control. The small orange sphere has fine curved ridges to indicate grip. Drag it along the time direction to seek and hear the recording without rotating the view. If playback was paused, releasing the sphere pauses at the chosen source time; if playback was running, it continues. Focus the sphere and use left/right arrows to seek or Space to play/pause. The ridged sphere sits at the top outer corner of the time plane nearest the viewer; the corner follows the viewing direction.

Octave third-octave and twelfth-octave bands

The Spectrum FFT button cycles FFT, Octave, 1/3 octave and 1/12 octave. Bands group FFT bin energy with window-energy correction. These are FFT estimates, not certified IEC octave-filter measurements. Bands narrower than the FFT resolution are marked unresolved. Labels are spaced to avoid overlap; the expandable Band levels list contains every band. Frequency weighting and the saved or estimated reference apply; source audio and replay stay unchanged. In octave mode the normal FFT curve stays visible above the translucent bars, with the same frequency and level axes.

Export CSV measurement data

In LYZ3R open Save / Export, choose Export to share, choose Tables only under Export as, then select the required Measurement tables (CSV): levels over time, frequency levels or octave bands. To include tables beside audio or video, choose that media output and expand the optional table section. The From / To controls set the exported interval. Level over time uses 100 ms blocks (the final block may be shorter), with RMS, sample-peak levels and sample counts; these are not Fast/Slow SLM readings. FFT exports the selected frequency bounds and current FFT size/window as bin levels, not power spectral density. Octave CSV uses the selected division, or one-third octave when the display is in FFT mode. Bands are base-2 FFT energy estimates, not certified octave filters; unresolved bands have an empty level and resolved=false. Files include source name, channel, weighting, reference status, level unit, sample rate, FFT settings, interval and reference gain. Estimated acoustic values remain estimates. Digital FFT bins use peak amplitude; acoustic FFT bins and octave bands use RMS. A value of -Infinity means exact zero signal. CSV uses commas, decimal points, UTF-8 and quoted text; use the spreadsheet import dialog if your locale expects semicolons. Original WAV audio stays unchanged. Save complete dataset keeps original media and metadata for reprocessing; CSV is a separate derived export.

08 LYZ3R zoom and display settings

Select the event first, then choose the frequency and level detail needed to see it.

Mark an interval in Channel

Channel is the master timeline. Drag across the waveform with one finger, or the mouse, to mark an interval in either direction. Drag the Start and End grips to refine it, or tap the labelled Start/End values to enter exact times. Spectrum, Spectrogram, Waterfall, 3CH0, replay and exports use that same sample-aligned interval. The source waveform keeps its context instead of jumping to the selected slice. Two fingers zoom or pan the visible waveform without changing the marked interval. Fit view shows the whole waveform without clearing the selection; Full recording clears the selection. A tap alone does not create an empty selection. The shared status readout shows interval start/end/duration, Average (RMS) and sample Peak in the current dB unit and weighting, independently for actual source channels. Marking never cuts or changes original audio. The rounded centre grip above a marked interval moves the whole selection without changing its length. The same highlighted interval appears in the optional full-recording overview. Full overview in Channel is off initially. Tap it to show or hide the selectable full-recording overview, including while zoomed. The same preference remains in View. You can mark in that overview too: the main Channel waveform and all connected results use the same interval. Small orange triangle grips sit above both marked edges, with a nearby rounded three-line centre grip for moving the interval. Their touch targets remain large. Channel marking is always available directly, without a Mark button. Multiple plots start compact. Drag the ridged grip at the bottom of a plot to change its height; keyboard arrows adjust it and Home or double tap restores automatic size. Resizing never changes the source or selected interval. Tap the editable colour-scale range to set both the dB span and colour palette. Editable endpoint values replace underlying printed labels rather than overlaying them. The same integrated interval RMS and sample Peak are also shown directly beneath the Channel Start and End time controls. They update when you mark, move or adjust the interval and do not use audio outside it. Channel marks time directly without a Mark button. The other analysis plots have a Mark option: Spectrum marks a frequency band; Spectrogram marks time and frequency together. Waterfall Mark uses its projected floor plane rather than rotating the graph. Turn Mark off to resume normal navigation. Time selections update the Channel origin and all linked views. Frequency marks highlight the same band and show Band RMS (FFT) and the peak-bin frequency at the bottom of every plot. Clear band removes only the frequency mark. Band RMS is an FFT estimate, integrated by summing one-sided bin powers with the selected window bandwidth correction, over the current time interval. It is distinct from the exact time-domain interval RMS. Sparse FFT frame coverage and finite frequency resolution limit the estimate, especially for short events or bands narrower than a bin. Marking never band-pass filters replay or alters the saved audio. The selected time and band are recorded in export metadata. Turning Mark off in Spectrum clears the shared frequency box and band values; the selected time interval remains. The other plots retain their existing marking controls. The smaller overall-peak label refers to the full visible spectrum over the selected time, rather than the marked frequency band. The boundary triangles point outwards with their flat edges aligned to the marked boundaries. The small rounded three-line centre grip moves the complete interval. Three wider, slim horizontal lines below each plot resize its height. Channel has a wider right-edge scroll area; dragging there scrolls the page without changing the selection.

Teal frequency endpoints and Log or Linear

Tap the upper/lower frequency values to edit them, for example 8k for 8 kHz. Pinch vertically to zoom frequency. Once frequency-zoomed, drag inside the spectrogram up/down to move the frequency band; drag diagonally to move frequency and time together when time is zoomed too. This works with other plots active. Drag the marked grip along the right frame edge to scroll between stacked plots without changing the analysis. Unzoom restores the full view. Log spaces frequency ratios; Linear spaces hertz. Original audio stays unchanged.

Top, bottom and Range

The colour-scale top and bottom are level endpoints. Type a top or bottom value to keep that endpoint fixed. Type a middle Range value, such as 30 dB, to follow the peak with a fixed 30 dB span. View also offers Lock colour limits for comparison and Auto range. Use fixed endpoints, for example top 80 and bottom 30, when comparing recordings on the same reference.

View, FFT size and Window

View contains FFT size (1024 to 8192), window type, palette, replay speed and frame controls. A larger FFT separates nearby frequencies better but spreads events over more time. Hann is a general starting point; Hamming, Blackman and Rectangular change leakage and detail. Changes recalculate analysis; they do not change source samples. The selected colour palette is shared by SLM, REC0R and LYZ3R, including their video data overlays. Change it in LYZ3R View; all three retain the same choice. Spectrogram and Waterfall also have a Dynamic button: enter a dB span to follow the peak and reveal quieter detail. Editing either level endpoint switches back to a fixed bound.

Can a long overview miss a short event

Yes. At most 180 analysis frames are distributed over the selected interval. Long overviews can leave gaps between frames, so zoom to a short interval around a clap or chirp before drawing conclusions. Spectrogram time labels refer to frame starts. This is narrowband FFT analysis, not an octave-band or power-spectral-density instrument. In octave mode the normal FFT curve stays visible above the translucent bars, with the same frequency and level axes. Native stereo PCM takes can be up to 64 MiB encoded so a five-minute 48 kHz stereo recording can be opened. Decoded and weighted audio uses more memory; very long stereo takes remain a phone-memory qualification limit.

Saved reference and estimated values

LYZ3R uses a recording's saved sensitivity/reference when available. Otherwise it labels an estimate using the generic phone assumption. Settings offers digital dBFS for input diagnostics. Keep the reference, weighting and colour endpoints consistent when comparing different files. A reference record does not certify the phone measurement chain.

09 Photos and video with measurements

Camera stays off until you activate it. Opening it always starts in Photo mode.

Photo mode and the shutter

The photo shutter takes a scene image with two seconds of original microphone audio and opens Save package. In REC0R, the red Record button still makes a normal audio take while Photo mode is selected. Close the camera to return to the ordinary plots.

Switch to Video and record

Tap the Photo/Video mode control to select Video explicitly. In REC0R the red Record/Stop button then controls video with audio. In SLM video follows Start/Pause; selecting Video while measurement is running starts video immediately. There is no separate second video-record button. Camera video is limited to one minute; SLM measurement may continue afterward.

Select Channel, Spectrum and Spectrogram

Activate Channel, Spectrum and Spectrogram independently in video mode. A white Channel envelope can overlay Spectrogram, while Spectrum uses a separate strip. Waterfall is for normal REC0R recording and LYZ3R analysis; it is not offered in video mode. Video-with-data replay/export reconstructs the supported Channel and Spectrogram overlays, rather than every analysis panel. LYZ3R export offers Channel, Spectrum and Spectrogram independently. Channel moves through a two-second source-time window; Spectrum follows the current source frame. The white spectrogram time marker continues moving. Video status can be Off, 1 line or 2 lines, including source time, units, weighting, channel, sample rate and FFT settings. Channel and Spectrogram share a transparent overlay, matching the live camera composition. Optional Spectrum uses the live teal trace. The white source-time marker remains visible. Short clips are analyzed within their exported interval using overlapping FFT windows.

Camera only and Video with measurements

Video with measurements keeps the scene, values and supported analysis graphics. The original camera stream and WAV are retained as sources in the package for later work. Camera only is an optional copy without measurement plots. Full app-view capture contains MEQ content, not Android notifications or the navigation bar. Recovering an earlier video never replaces a newer photo or audio take; older video is linked only when its original WAV matches the current audio.

Camera stops and watermark

Capture stops when switching modules, backgrounding the app, losing its input, reaching one minute or approaching the buffer limit. Save the stopped take before another capture. The MEQ logo and timestamp identify the export; a watermark does not authenticate a recording. Device audio/video latency is not a calibrated delay correction. Android audio-only REC0R is the exception: with screen-off recording enabled, it continues until Stop or the low-storage guard, retaining completed four minute parts. Camera video, live preview and SLM still require the foreground.

10 Replay crop and save packages

A package preserves the sources needed for later analysis. A PNG is only an image.

Play, Stop, Loop and Speed

In Replay, choose Audio or Video with data when available. Play starts or pauses; Stop ends playback. The seek slider changes position. Loop repeats the From/To interval. Speed offers normal, half and quarter speed with pitch preservation where supported. Slow listening does not change the original measurement timeline. Measurement-video copies omit navigation and interactive controls; their scene, selected plots and optional status use the export preview. In REC0R, Replay options contains range, speed and plot choices. Restarting, looping or seeking backwards resets the spectrogram history so it follows the new playback position.

From, To, Full take and crop export

Drag the time handles or type From/To seconds. Full take restores the complete range. Export creates a selected-range copy; Save complete dataset retains full originals. Apply replay speed to exported copy is a separate Replay opt-in. Listen to export prepares the actual audio; Play export prepares video; Preview table shows the selected table. Export reuses that prepared file while settings stay unchanged. Changing range, including Full take, audio routing or video appearance clears an obsolete preview. Moving Look at time only inspects appearance and keeps the prepared file. Controls are locked during preparation/saving. Video export processes the selected samples and frames locally instead of waiting for real-time playback. Clips can be prepared faster than their playback duration, depending on the phone, codec, clip length and selected plots; the first export can take longer to initialize. Normal, half and quarter speed use the selected source interval; slowing preserves pitch in the copied audio. The rendered copy is limited to one minute. Accurate video encoding requires supported codecs and an up-to-date Android System WebView or browser. If decoding or encoding fails or stalls, preparation reports an error and retains the originals. Video with measurements also works with an audio-only take; no camera recording is needed. Choose Channel, Spectrum, Spectrogram and Waterfall independently. Use visible views copies the current analysis view choices; Select all includes all four. Without camera video the selected views appear as separate panels with original-source time and the take audio. Waterfall axis controls use transparent backgrounds so they do not cover the surface. The export preview and video keep the Waterfall rotation, tilt and viewing proportions last set in analysis; its surface is prepared once while the time slice continues to move. Camera video retains its transparent overlay composition; Waterfall is an audio-only panel option. Preview export prepares the actual selected output. Saving with unchanged settings reuses that prepared file. Play export prepares and starts the actual finished video. If Android requires another user gesture, tap Play export again or the player Play control; the prepared file is reused. Inspect one frame is optional and closed initially. Its Look at time slider only inspects appearance and does not freeze the exported movie. Audio-only video export also supports takes opened directly with REC0R Analyze, as well as imported recordings. The export range also has a bar-waveform view with draggable start and end grips. The greyed areas are outside the selected export interval. Drag either grip or use From and To; both channels always use the same sample interval. Keyboard arrows adjust the focused grip, with Shift for larger steps. The original source remains unchanged. Save complete dataset keeps the full original; Export to share creates the trimmed copies. The marked interval is exported directly. Keep marked retains the inside. Remove marked joins the outside audio and is offered for WAV audio only. Full take resets the interval. Preview prepares the exact file that will be saved. Both stereo channels use identical boundaries; original files remain unchanged.

WAV, video, photo, JSON and TXT

Save complete dataset keeps full original WAV and camera media, available documentation copies, notes, annotations, source reference and structured take.json metadata. WAV is the original sound for listening and reanalysis. TXT notes are an optional readable companion. Keep the package files together when downloading separately in a browser or saving without batch support. A plot PNG alone does not preserve these sources. Export to share instead makes finished copies: choose Audio (WAV), Video with measurements, Camera video only, Photo where available, or Tables only under Export as. The same choice controls the preview and exported file. Measurement tables (CSV) can also accompany media through their optional collapsed section. Confirmed saves list written files and destination; Back becomes Done. Cancellation or failure does not show a completed success state. For stereo audio, a complete dataset includes the original stereo WAV and aligned Left and Right mono WAVs, with channel names and sample counts in the JSON. A WAV export also produces Left and Right copies from that same prepared, trimmed interval. Stereo channel separation is preserved; the mono-to-both-ears option does not duplicate a silent stereo channel.

Where do all files go

Choose recording folder in Settings selects a shared folder once. Documents / MEQ App is suggested; an existing chosen folder is retained. Settings shows the folder and offers Open recording folder when a compatible viewer is available. Save shows its destination and preparing/writing progress. Wait for completion before sharing. If some export files save before an error or cancellation, the dialog reports how many succeeded and keeps Share saved files available. Retry remaining writes only outstanding files with unchanged settings. Changing output, range or other export settings starts a new attempt; earlier saved files remain in their folder. Original recordings stay in the app.

Share saved files and WhatsApp

Share export opens Android sharing with every file just exported, including selected CSV tables. A partial export can share its successfully written files. Preview saved file opens the chosen result in a compatible viewer. Complete-dataset sharing offers a video, photo or audio copy while the full package remains in its folder. MP4 support depends on the phone codecs and receiving app; WebM is an alternative where offered. Keep original media and take.json even when sending a smaller compatible copy.

How does importing find the right video

A saved package carries its file association. For separate audio, automatic video linking requires an explicit package identity or one unique exact audio-file match. Similar names, sizes, capture times or durations are suggestions, not proof. Ambiguous matches stay unlinked. Use Change / add video for a manual selection and verify the scene and timing. An unverified manually attached video does not replace the audio source reference, calibration information or notes. Original samples remain unchanged.

Recognizable date-first package names

New saved package files start with a timestamp, followed by the optional short Take title and file kind. The original capture timestamp is used when known; otherwise the name uses export time and metadata leaves capture time unknown. Existing files are not renamed. Package metadata retains the original source and reference. Enter a short title such as Elephant clap to make later filename and note searches useful. Only processed A/C acoustic plot, CSV and measurement-video copies receive _dBA or _dBC. Raw WAV and original camera files remain unweighted and never receive those suffixes. Exported WAV is raw audio, optionally cropped; choosing weighted plots does not filter or overwrite it.

Readable video text

Video export preserves the working live composition: a transparent spectrogram with the thicker white Channel trace on the same time axis, an optional Spectrum, and the moving white source-time marker. Quiet regions reveal the camera scene; stronger signals have stronger colour. It uses the appâ€™s Outfit font and adapts to video proportions without separate titled graph panels. The first LYZ3R export follows the selected analysis plots; later export choices are retained. The appearance preview updates with plot/status choices. Look at time scrubs the scene and linked marker; audio shows a time view and tables show their first rows. Small status is optional: Off, 1 line or 2 lines. The appearance preview is a still; Listen to export or Play export prepares the actual finished file, including sound. Camera video only adds no measurement overlay. Mono / left-only audio to both ears is an export-only option, on by default. It duplicates mono or a left channel with a completely silent right channel. Genuine stereo stays stereo. Normal replay and complete-dataset saving keep the original channels. Display formatting does not change measurement values or original samples.

11 GEN3R8 sources and output

Generate a tone/noise on each side, or combine an audio file with a generated signal. Start with low device volume.

Sine, White, Pink and Brown

Sine is a single-frequency tone; type its frequency in hertz. White, Pink and Brown are different noise spectra. Level controls the digital source amplitude/ceiling, not an acoustic dB SPL target. The selected output device and phone volume still affect what reaches a loudspeaker.

Link and Split channels

Linked channels share source settings, copying Left to Right when linked. Split lets Left and Right have different waveforms, frequencies, levels or files. Output trims and mutes remain independent. Changing the link stops playback; press Play again when ready. Stereo separation requires an output path that preserves left/right channels.

A song on one side and noise on the other

Split channels. Load an audio file on one side and choose White noise, Sine or another source on the other. File sources loop; stereo files are downmixed to mono for that selected side. Loaded files clear on app reload and are not embedded in the synthetic generator preset/WAV export. Check the visible source names before starting.

Output L/R, trims, mutes and meters

Output L/R adjusts each output trim and mute after source settings. Digital meters are before device volume, so a low phone volume does not lower their displayed digital level. Software Stop ends playback; it is not a hardware emergency stop. Digital values are not calibrated volts or sound pressure. The marked interval is exported directly. Keep marked retains the inside. Remove marked joins the outside audio and is offered for WAV audio only. Full take resets the interval. Preview prepares the exact file that will be saved. Both stereo channels use identical boundaries; original files remain unchanged.

Save Load and portable presets

The generator file sheet exports a stereo WAV for synthetic sources and saves/loads a portable JSON preset. WAV is playable sound; the preset is editable settings for a later session. File-based sources do not support the synthetic-only export/preset route. The drawing/sweep Files sheet is separate and saves its own editable settings.

GEN3R8 Replay

Replay loads supported audio, with independent source gains and output trims. Typical supported inputs include WAV, MP3, MP4 and M4A when the audio track can be decoded. Limits are 100 MB encoded, ten minutes, mono/stereo and 64 MiB retained decoded audio. LYZ3R is the module for waveform/spectral inspection and measurement-video work. If device audio is interrupted, Replay pauses at its current position. Press Play explicitly to continue; returning audio focus alone does not restart playback.

Musical note presets

In Generate, the Notes button beside each frequency slider opens the standard 88-key piano range A0 to C8. All chromatic semitones are included, with sharps and flats shown together. The reference is A4 = 440 Hz in twelve-tone equal temperament. Filter by octave or scroll, then choose a note. This selects a sine tone for that channel, stops current playback and sets the frequency without starting sound. Press Play to listen. When Left/Right are linked, the Left control follows the existing linking behaviour; split channels for independent notes. The original numeric frequency and slider remain available. These are reference tones, not an implemented microphone tuner. The larger Musical notes button, with a musical symbol, is immediately after Brown in the source choices. Import audio, with a folder icon, is a separate control for existing recordings.

12 GEN3R8 Sweeps and Draw

Sweep makes a moving tone. Draw turns a time-frequency drawing into sound.

Start End Time Curve and Direction

In Sweeps, set Start/End frequencies and duration. Linear changes frequency by equal hertz per second; Log changes by equal frequency ratios. Direction chooses the sweep direction. Level sets the digital ceiling and Loop repeats playback. Sweep output is a mono source sent to both channels.

Draw and Play without Render

Select Draw and draw in the large plot with a finger, pen or mouse. Horizontal position is time over ten seconds; vertical position is frequency from 0 to 10 kHz. A horizontal stroke produces a steady sine tone, while a sloped stroke changes frequency. Press Play to create the sound automatically; the separate Render step has been removed. Prepared sound is reused until you edit it.

Brush Undo and Clear

Brush opens Visual brush width, strength, output level and loop settings. Width changes only the thickness of new strokes on the drawing. Each stroke plays one continuous tone following its centre frequency; strength changes its sound contribution. Undo removes the latest stroke; Clear removes the drawing. Separate strokes add together. The fixed Jet drawing colours and visual gain help you see strokes; they are not calibrated acoustic levels.

Files Save settings Open settings and Export WAV

Save drawing / sweep settings stores editable strokes and configuration in a JSON project. Open drawing / sweep settings restores such a project; it does not import an ordinary song or WAV into the drawing. Export WAV saves the generated sound. Press Play once after editing to prepare the sound before exporting. Keep the project too if you want to revise it later.

13 SL33P scenes and timer

SL33P is its own module at the end of the default app menu.

Scenes White Pink Brown and Mix

Choose a named scene or a direct White/Pink/Brown noise shortcut. Scenes expands the available choices inline. Mix adjusts the contributing layers. Generated scene names describe an inspiration, not an authentic location recording. Adjust listening level and press Play.

My audio and named recordings

My audio adds a local recording and offers its volume, Loop, Recording only and Remove. Named Rain/Jungle recordings use a generated fallback when the corresponding file is unavailable. Local selections clear after reload. These recordings are not automatically included inside a saved preset.

Sleep timer

On Android, native SL33P playback continues when the app is hidden or the phone locks. Timed playback uses a native timer with a final ten second fade. Until stopped keeps playing until Stop or an audio interruption. The playback notification has Stop. Headphone removal and loss of audio focus stop playback; press Play explicitly to resume. Switching modules stops the previous module. Generated sound uses a seamless playback loop; imported sound plays separately with its chosen loop setting. Mix recording is foreground-only and finishes when the app is hidden. Browser-only playback does not provide this native background service.

Record mix and save WAV

Record mix captures the playing mix locally as PCM WAV, up to five minutes within a 32 MiB capture budget. Save WAV keeps the result. MP3 export is not included. Recording a mix and setting the sleep timer are separate actions. Mix recording is foreground-only. Hiding the app finishes and retains the current WAV mix recording while native soundscape playback can continue.

14 Take information GPS and environment

Add information that makes a recording understandable later, then save its package to keep it.

Info Take name Location and Notes

LYZ3R Info and Save package / Info & notes open the take information sheet. Give the take a meaningful name, location/setup and observations such as distance, microphone position, wind and what happened. Editing these fields is not a file save; Save package writes them into the take metadata.

Add current GPS

This explicitly requests the phone's current location. The coordinates, accuracy and their own fix/added timestamps are recorded. Adding GPS while inspecting yesterday's take does not reconstruct yesterday's capture location. Review what you share: a saved package can contain precise coordinates.

Temperature Humidity and nearby weather

You can enter temperature and humidity manually. Before an optional nearby-weather request, a confirmation explains that coordinates rounded to two decimal places go directly to MET Norway, which also receives your IP address. Cancel sends no weather request. The saved model temperature, humidity, wind, source and valid/retrieval times are nearby estimates, not phone measurements or a reconstruction of past capture weather.

Recording and device information

Recording & device information uses readable labelled values: original device/source, capture time and time zone, capture app version, duration, sample rate, channels, audio format and reference. Viewing on this device is separate. Not saved means the original value is unavailable. Advanced diagnostic data retains full metadata. New audio-only recordings retain capture device/time-zone/app-version provenance too. Notes and the shared recording Save dialog remain optional; no extra reminder appears after each take.

Optional dictated notes

Dictate notes is available beside Notes in both Info and the existing Save dialog. It uses Android on-device speech recognition when available; it never silently switches to an online service. Stop recording, measurement and Live View first. Tap Dictate, speak, then Finish or pause. Recognized text is appended for review and editing; existing notes remain. Closing the dialog cancels dictation. Device language/model support varies; typing remains available. Notes stay optional and do not trigger automatic popups.

15 Frequently asked questions

Short answers to the questions that shaped the app.

What is dBFS Why are the numbers negative

dBFS is a level relative to digital full scale. In this app, an amplitude of 1 is the reference; quieter digital signals are below 0 dBFS. It is useful for checking recording headroom, but is not a sound-pressure reading. Measurement mode shows estimated dB by default; the digital convention is an expert setting.

What is the difference between dB dBA and Z

dB needs a reference. Here the display distinguishes estimated acoustic values, a saved reference and digital dBFS. dBA indicates A frequency weighting. Z is the unweighted option; choose it when you want no A/B/C curve. Changing weighting changes the displayed contribution of frequencies, not the original WAV.

Must I calibrate before I can start

No. Start immediately with estimated readings. Optional reference adjustment improves correspondence to the chosen reference under that setup, but does not qualify the microphone, phone processing or entire meter. Do not mistake Estimated for a certified result.

Why is a save or video button grey

The corresponding output is not available yet, capture may still be active, or a matching original video has not been attached. Stop the take first. A sound-only WAV cannot supply scene video. A source video plus its WAV/package enables later video-with-data work.

Can I remove a baked in overlay from an old video

No. An overlay already encoded into video pixels cannot be recovered as a clean camera original. You can still use that video as the scene reference and analyze its separate WAV. New packages retain the original camera stream for this reason.

Why is the waveform large or the colour suddenly brighter

The waveform can be automatically scaled, and Dynamic can move the spectrogram colour range. These are separate display scales. The chosen 60 dB setting means the spectrogram colour endpoints differ by 60 dB; it does not set the white Channel amplitude or Spectrum scale. Read the units, endpoints and reference status before comparing recordings. Original samples are unchanged.

Frequently asked questions continued

What did Quiet to loud mean

It described the colour direction. The current plots use numerical level limits and Peak/Low or RMS information instead, so you can make a meaningful comparison. Use the editable colour endpoints to set a common range.

What are digital input diagnostics

They show the recording's digital signal level and input details for troubleshooting. They are optional expert information, not an extra acoustic measurement or a required startup step.

Why did my recording stop or disappear from the workspace

Switching modules, backgrounding, locking or a capture limit can stop recording. A new capture or Clear changes the current workspace. Packages already saved remain in the selected phone folder. Re-import the package; retain all original files if recovering an interrupted recording. Android audio-only REC0R is the exception: with screen-off recording enabled, it continues until Stop or the low-storage guard, retaining completed four minute parts. Camera video, live preview and SLM still require the foreground.

Does PNG save my WAV or video

No. PNG saves a plotted image with its units/metadata. Save package keeps source audio, available media and take information. Both use the selected MEQ phone folder in the consolidated update. PNG opens an exact image preview first. Save PNG writes that preview; Back returns without saving.

Where should I start for a clap and its reflection

Record a short REC0R take and stop. Analyze it, select Channel plus Spectrogram, and squeeze the time handles around the event. Replay at half or quarter speed to listen; use original-time positions for delay interpretation. Show video frames if available, then crop/export a copy and save the originals with notes. Device capture latency and phone processing remain unqualified.

How do I scroll when a plot is zoomed

Use the marked vertical grip on the right edge of an analysis plot, or drag its header/frame. The inner spectrogram pans time and frequency when zoomed; the separate edge scrolls the page and leaves your axis settings unchanged. This also works when several plots are stacked.

Does Together combine my recorded microphones?

No. Channel and Spectrum draw both original channels in one plot. Together Spectrogram and Waterfall show the mean of their powers and are labelled L + R. Separate shows each channel independently. The saved original stereo source, Left WAV and Right WAV retain their samples and timing. Red OVERLOAD holds for three seconds after the latest near-full-scale digital peak; it cannot prove that an analog microphone or receiver is free from overload. The marked interval is exported directly. Keep marked retains the inside. Remove marked joins the outside audio and is offered for WAV audio only. Full take resets the interval. Preview prepares the exact file that will be saved. Both stereo channels use identical boundaries; original files remain unchanged.

Where is my recording and how do I delete it

Clear resets the workspace and keeps retained recordings. In Voice use Manage retained recordings, select entries and Delete selected. In SLM use Open and Delete next to a retained measurement. Confirm permanent deletion. This removes local originals, edits, transcript and measurement history. Exported files and older device backups must be removed separately. New builds exclude app data from Android backup and device transfer. Continuous capture keeps completed four minute parts separately. Use Open continuous recording parts to open or delete them; the final partial part is retained after Stop.

Cut out the middle or invert a selection

In Voice, drag on Channel to mark an interval. Keep marked retains the inside. Remove marked cuts out the interval and joins the remaining outside audio. Use Undo or Restore original if necessary. Zoom and pan do not change the original. In export, the marked interval is used directly; Remove marked exports the outside audio. Save complete dataset keeps the full original.

Microphone not working or wrong input

Allow microphone permission in Android settings. Stop other capture apps. In REC0R use the input check and confirm the actual routed device, channel count and rate. Reconnect USB before selecting it. SLM uses the device route reported by its browser capture; this is not proof of certified accuracy. Stop and recheck the reference after input or gain changes.

No wifi and offline functions

Audio recording, editing, generated sound, analysis and bundled Voice transcription work locally. Optional weather uses an online service after you request it. Translate hands selected text to an external app which controls its own connection. Exports and sharing send only the files or text you choose.

Keep playing overnight and screen locking

SL33P uses native media playback on Android, with a Stop notification and a native sleep timer including a final fade. Headphone removal or audio focus interruption stops playback. Return and press Play to resume. Generated layers are prepared as a seamless bounded playback loop; imported sound plays separately at its original duration with the chosen loop setting. Verify your device battery settings. SLM capture still stops safely on screen locking and retains completed blocks.

LCpeak noise dose and supported metrics

Current SLM readouts include time weighted level, integrated equivalent level and maximum time weighted level. A maximum Fast or Slow value is not true instantaneous LCpeak. Certified noise dose, occupational compliance, microphone response file correction and deconvolved RT60 are not provided. Use a qualified measurement system for regulated decisions. Replacing the microphone alone does not certify this app.

Analyze the audio you just edited

Stop capture. In Voice or SLM choose Analyze this audio and confirm replacing the current analysis view. The selected working audio opens in Channel and feeds its linked spectrum, spectrogram and waterfall. Saved originals remain separate. PCM WAV import is limited to 64 MiB; lower rate audio retains its original bandwidth even when SLM internally resamples it for weighting.

Mobile app privacy and retained data

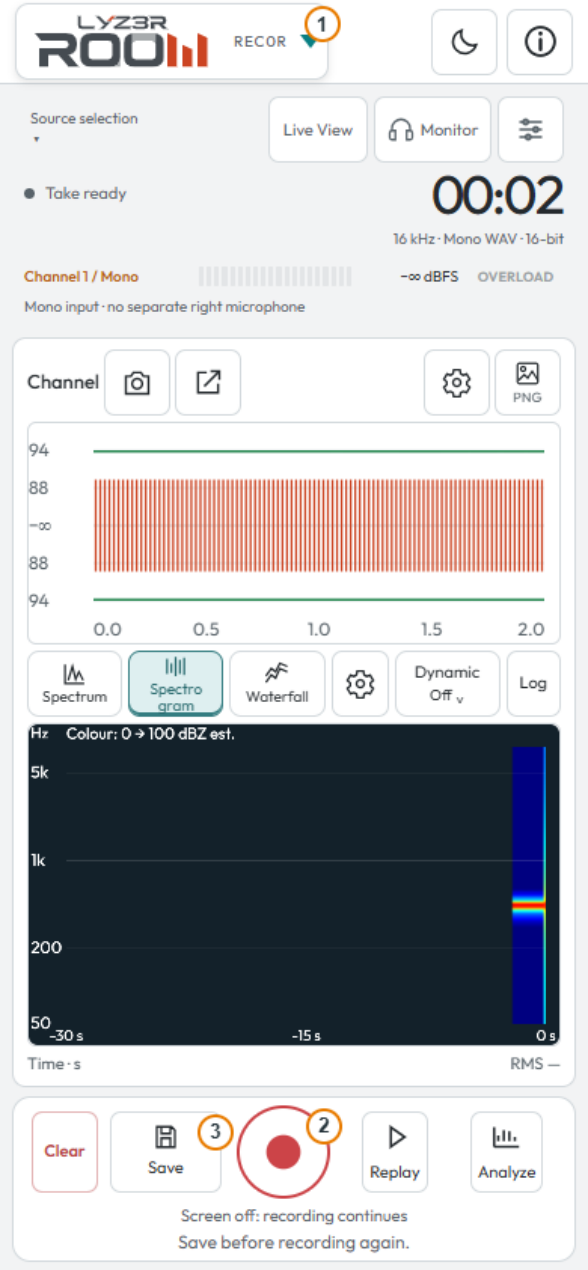
Recordings, edits, transcripts and measurements are retained on this device until explicitly deleted. Clear is not Delete. No account is required. The app excludes its data from Android backup and device transfer in this version; earlier backups may still exist. Optional location and weather are requested by you; weather sends rounded coordinates to the weather service. Translate sends chosen text to the external app. Sharing sends selected files, which may contain notes, timestamps, location and device metadata. Review metadata before sharing.

16 Illustrated quick reference

Numbered examples identify the main controls in each module and tab. Demonstration recordings, transcripts and camera scenes are synthetic; these pictures do not document a physical measurement. Use the matching task topics for detailed instructions.

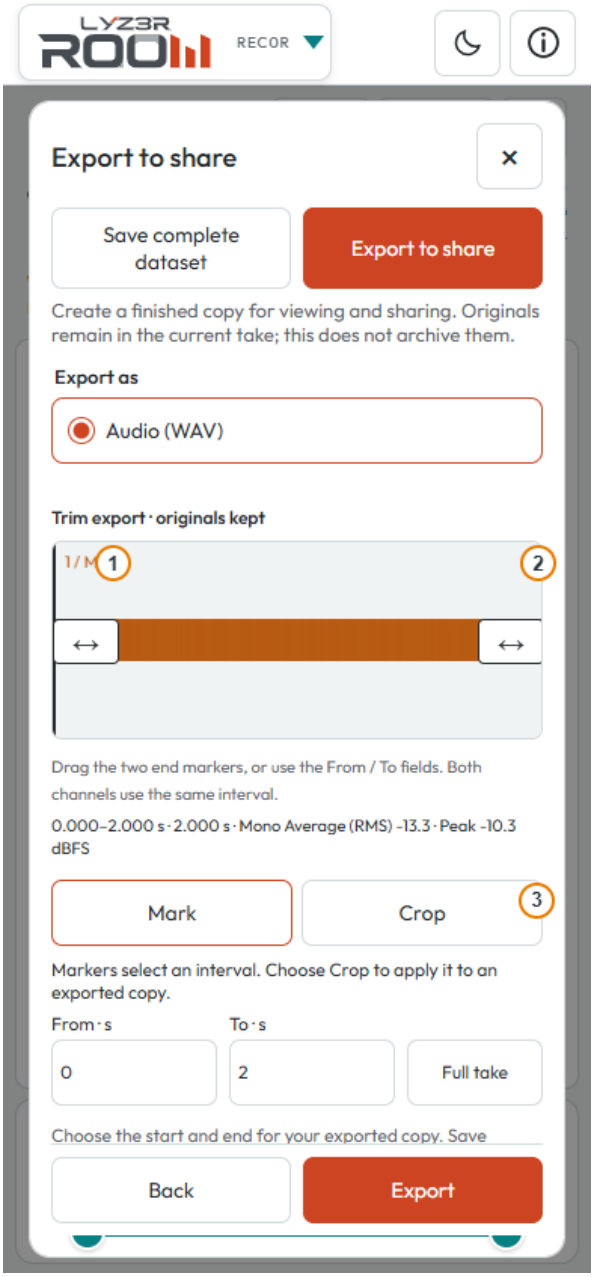
RECOR: record and save

1 Choose a module from the ROOM header. 2 Record / Stop. 3 Save full originals or export a finished copy. Source shown here is synthetic.



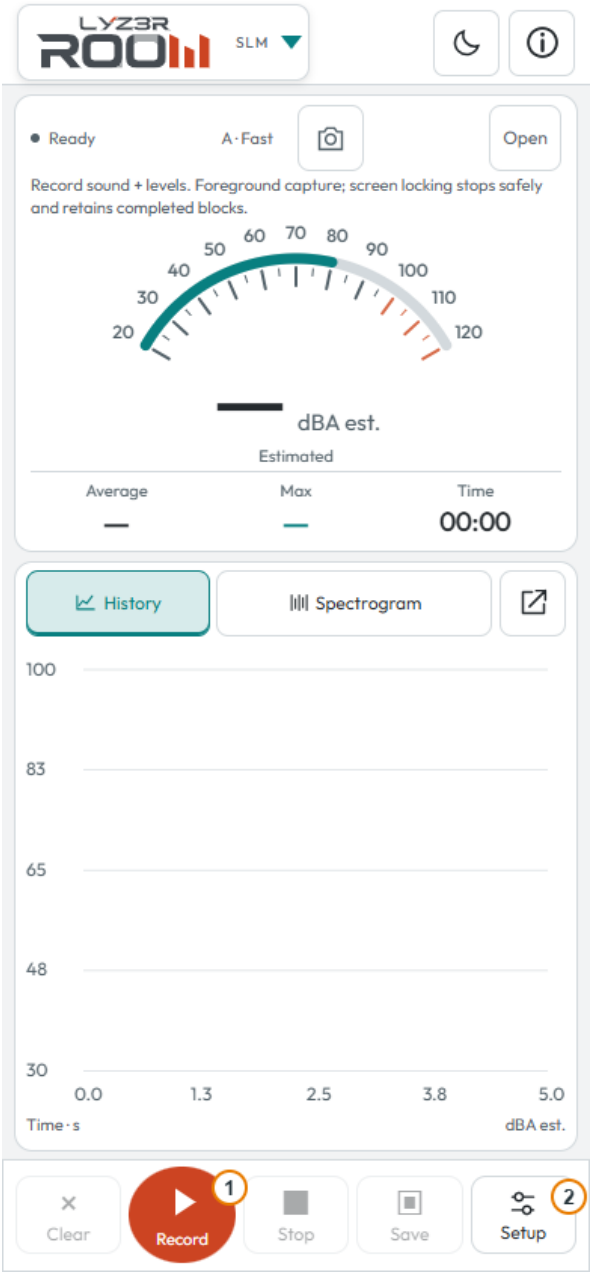
Mark, then crop

1 Start marker. 2 End marker. The level readout integrates between them. 3 Crop opens Keep selection or Remove selection; preview and save a copy.



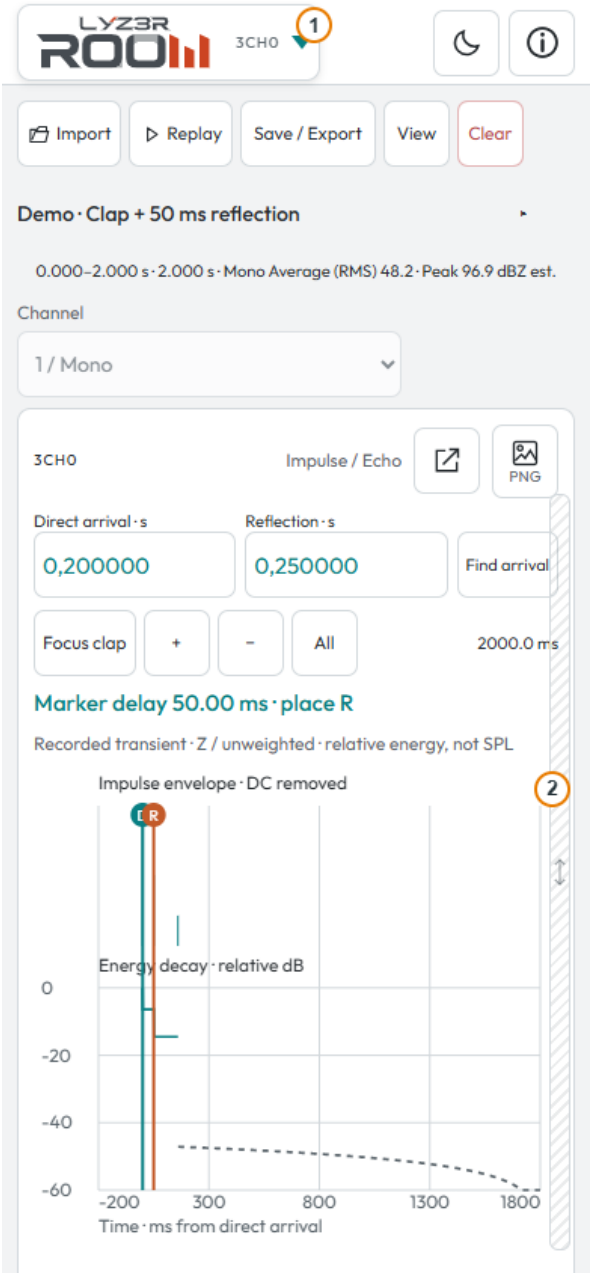
SLM: level and setup

1 Record / Pause; Stop completes the session. Save retains sound and measurements; Open restores saved sessions. 2 Input, weighting and reference setup. Levels support initial investigation, not certified measurement.



3CH0: reflections

1 3CH0 uses the shared analysis source. 2 Inspect impulse timing and reflections; this example is simulated. Drag the orange marker or zoom for detail.



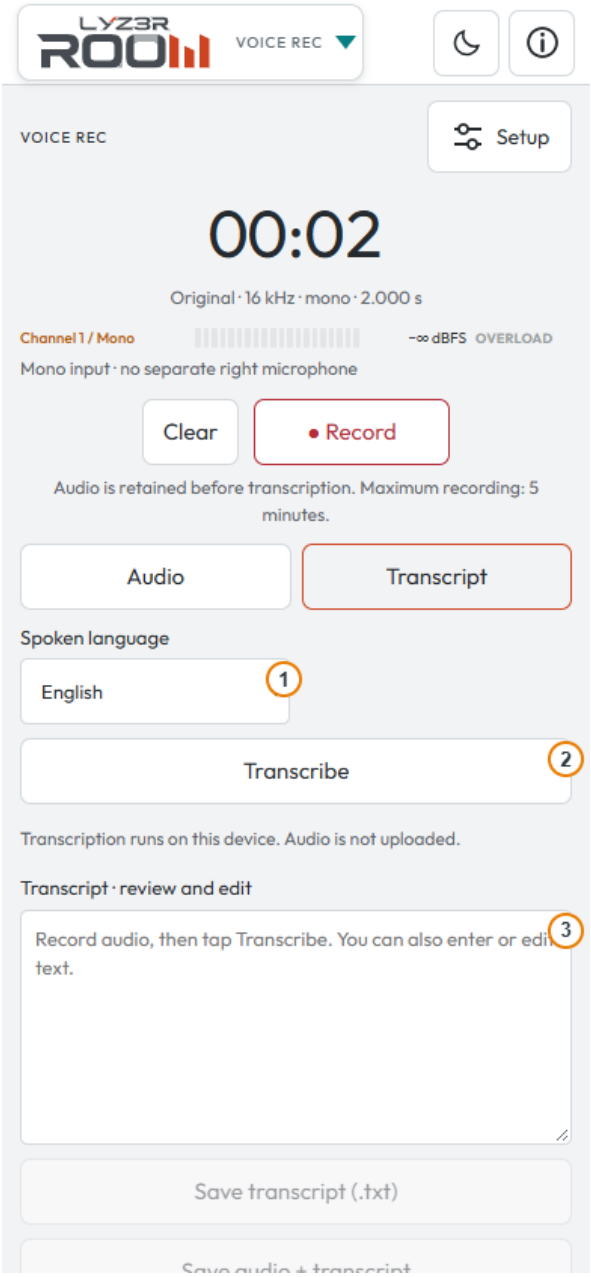
VOICE REC: audio

1 Record or Open audio. 2 Drag to mark and zoom as needed. 3 Keep marked retains the inside; Remove marked joins the outside. Undo restores the previous edit; the original is retained.



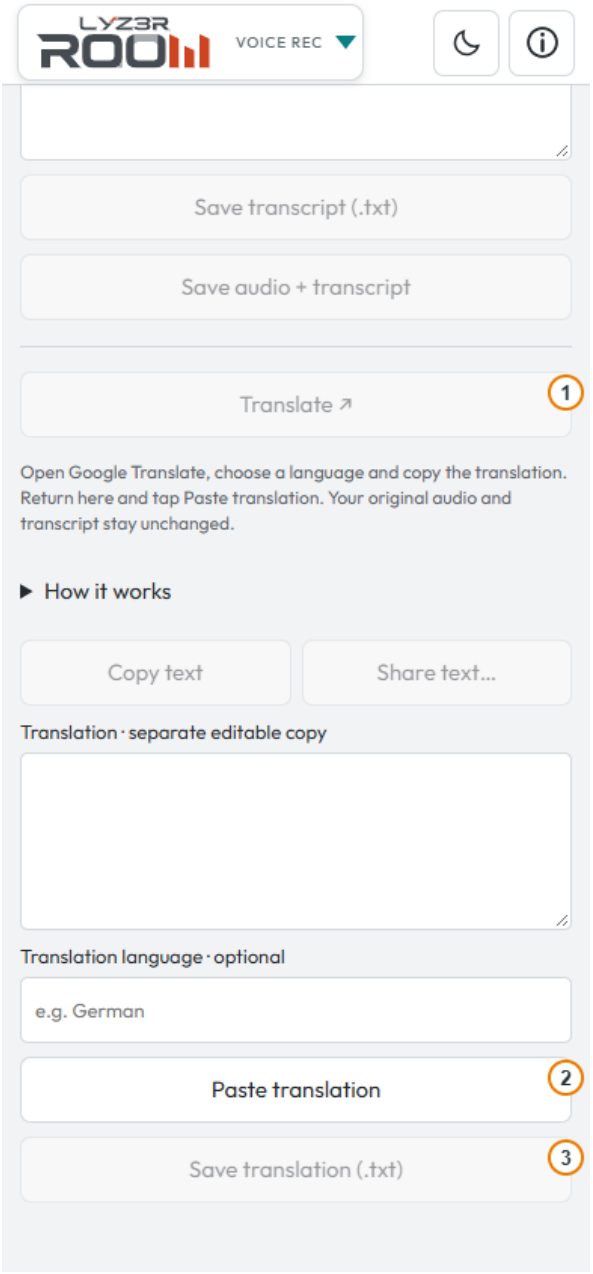
VOICE REC: transcript

1 Spoken language. 2 Transcribe locally on the device. 3 Review and edit the draft, then save the complete TXT. Recognition can make errors; compare with replay.



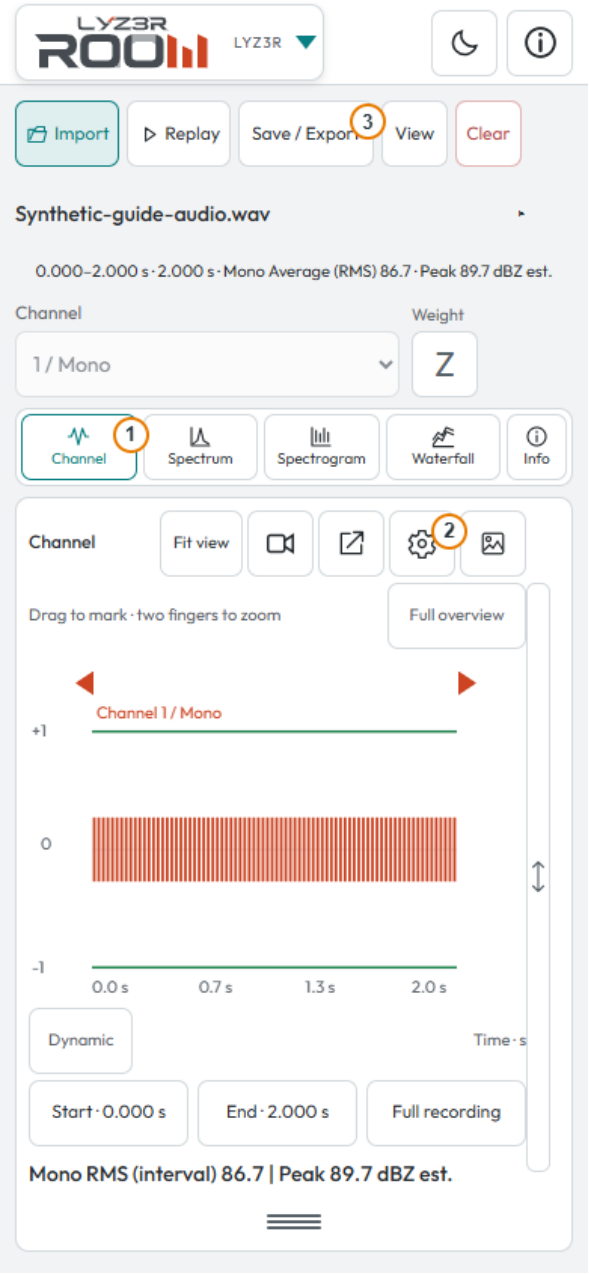
VOICE REC: external translation

1 Open Google Translate, choose the target language there and copy the result. 2 Return and explicitly paste into the separate translation field. 3 Save translated TXT. The original transcript remains unchanged.



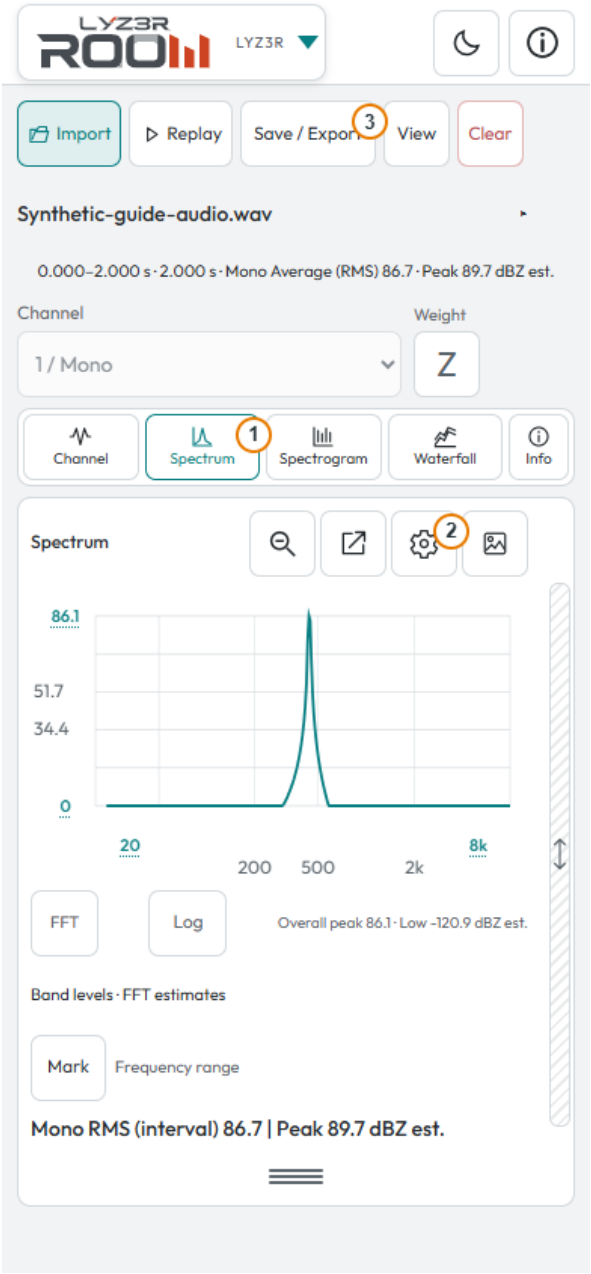
LYZ3R: wave

1 Select the analysis view. 2 Plot settings: together/separate channels and relevant display controls. 3 Export copies. The source is synthetic; selections and display adjustments do not overwrite original audio.



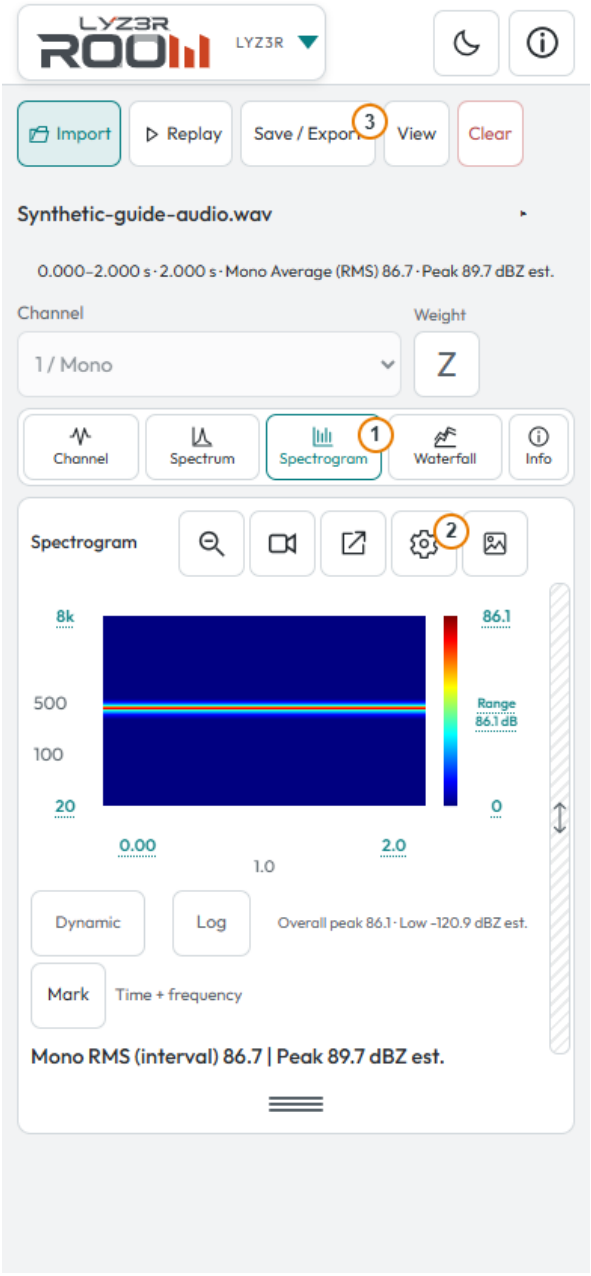
LYZ3R: spectrum

1 Select the analysis view. 2 Plot settings: together/separate channels and relevant display controls. 3 Export copies. The source is synthetic; selections and display adjustments do not overwrite original audio.



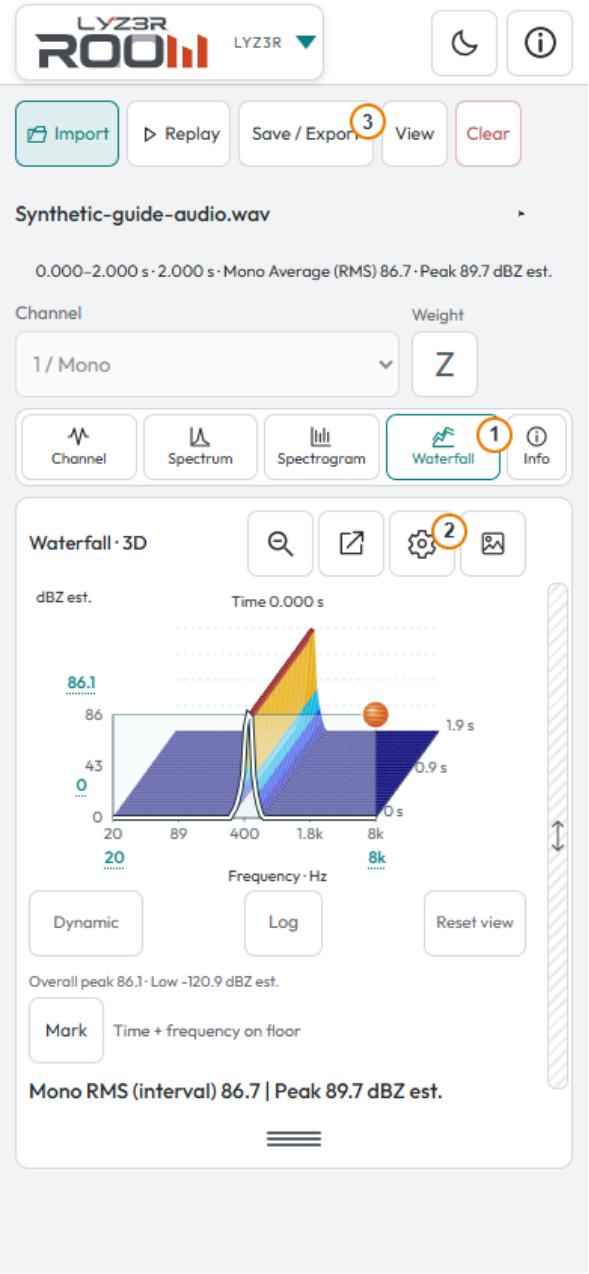
LYZ3R: spectrogram

1 Select the analysis view. 2 Plot settings: together/separate channels and relevant display controls. 3 Export copies. The source is synthetic; selections and display adjustments do not overwrite original audio.



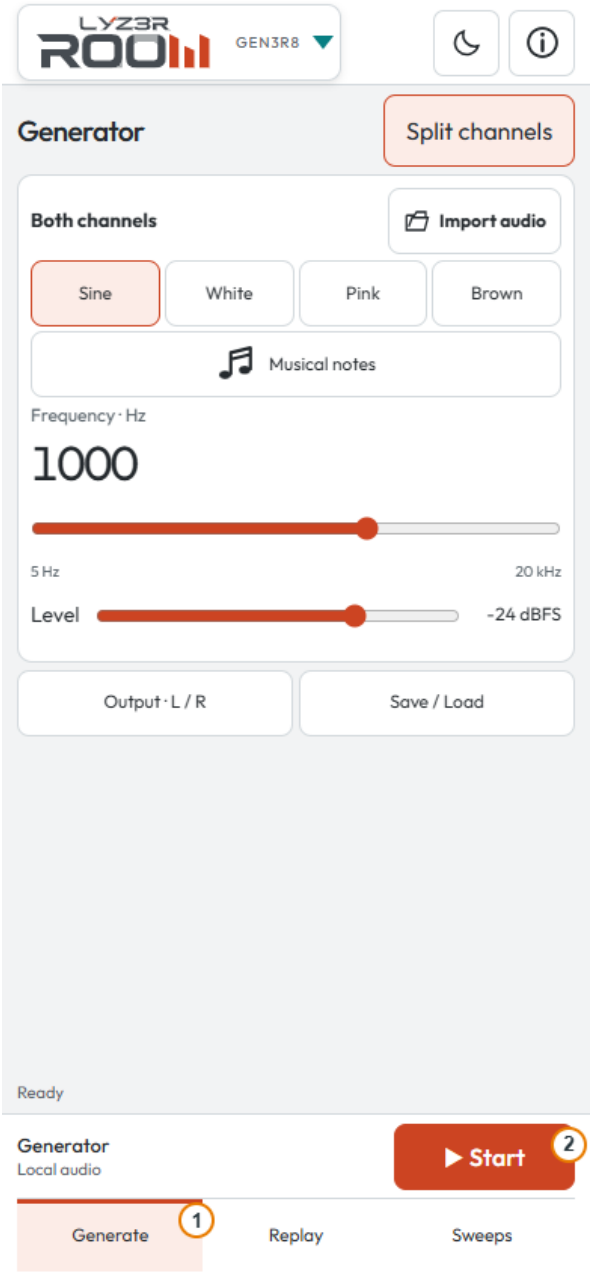
LYZ3R: waterfall

1 Select the analysis view. 2 Plot settings: together/separate channels and relevant display controls. 3 Export copies. The source is synthetic; selections and display adjustments do not overwrite original audio.



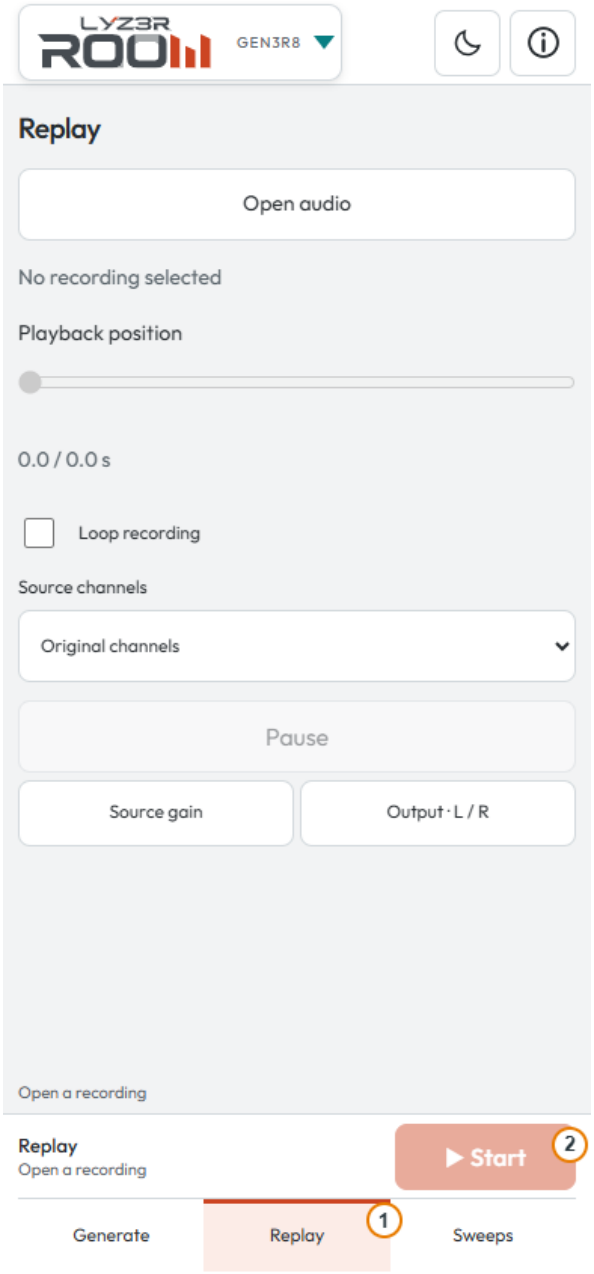
GEN3R8: generate

1 Select generate. 2 Play / Stop. Playback starts only when you request it; check output level before listening.



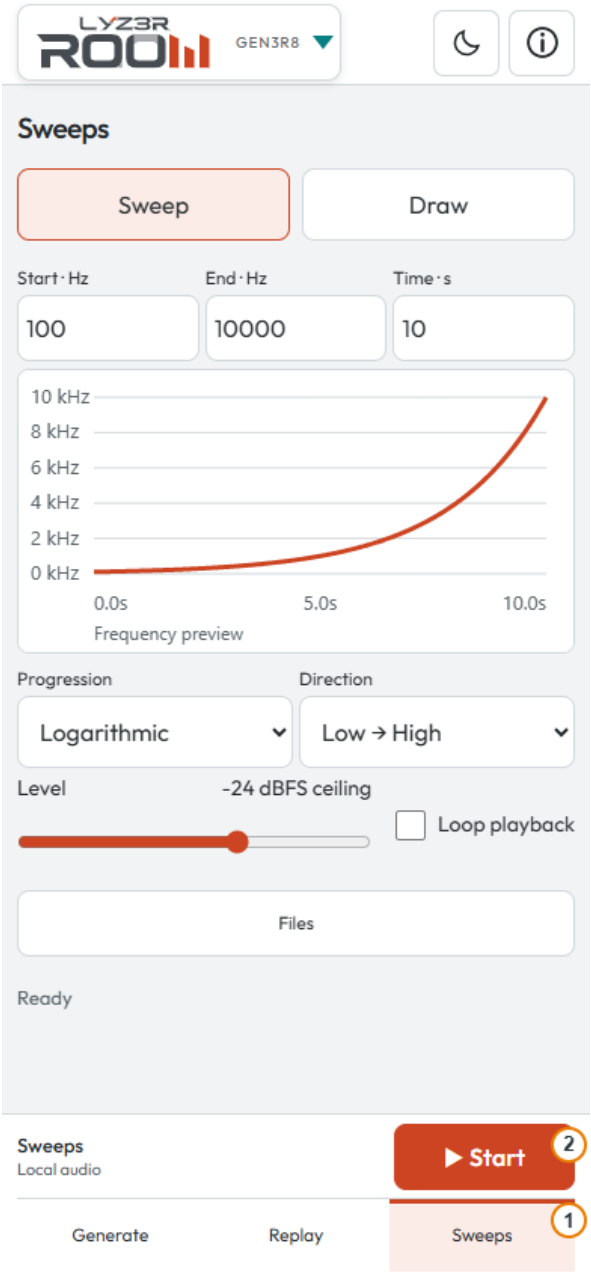
GEN3R8: replay

1 Select replay. 2 Play / Stop. Playback starts only when you request it; check output level before listening.



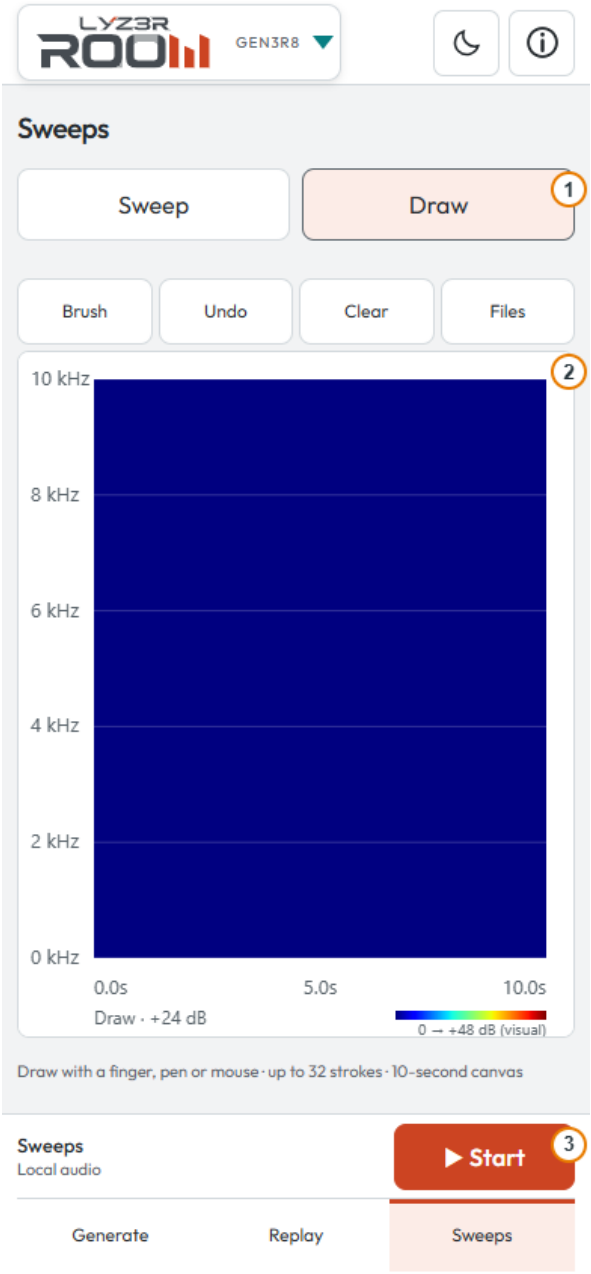
GEN3R8: sweeps

1 Select sweeps. 2 Play / Stop. Playback starts only when you request it; check output level before listening.



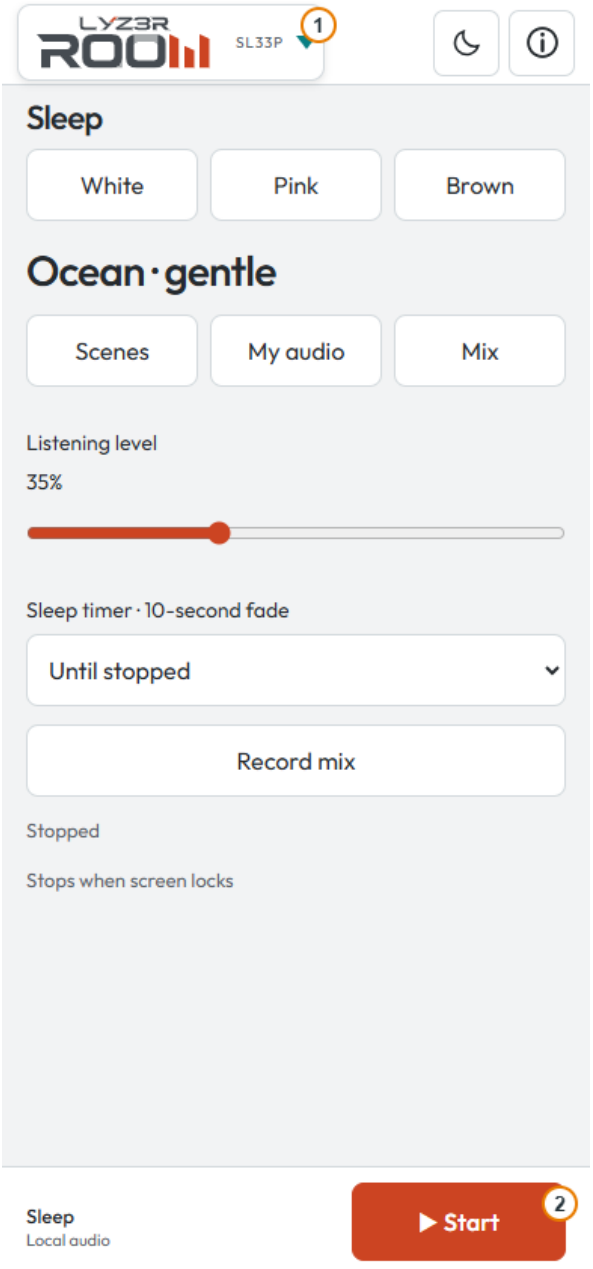
GEN3R8: Draw

1 Draw mode. 2 Paint the desired sweep strokes. 3 Play prepares and plays the drawn sound. Existing Draw colours and level behaviour are retained.



SL33P: sounds and timer

1 Choose SL33P. 2 Play / Stop the selected sound; use its timer controls for the desired duration.



Feedback to MEQ

1 Draft feedback locally. 2 Copy it. 3 Open the MEQ contact page and paste into its form. Opening a page or an email draft does not send feedback automatically.

