


fluffyAUDIO

FFT

SPECTRAL AUDIO EDITOR

User Manual

The complete guide to recording, editing, repairing and mastering audio on macOS and Windows — waveform & spectrogram editing, 47 processing modules, batch and on-device AI.



Version 0.2.61 · 47 modules · macOS & Windows

Contents

Getting Started

- Welcome to FFT
- The FFT Window
- Opening, Importing & Recording Audio

Playback, Navigation & Editing

- Playback & the Transport Bar
- Navigating & Zooming
- Selecting Audio
- Core Editing
- Undo & Edit History

The Spectrogram & Spectral Editing

- The Spectrogram & Spectral Editing

The Module Reference

- How Modules Work
- The Module Chain

REPAIR MODULES

- Spectral Repair
- Interpolate
- Spectral De-noise
- De-reverb
- De-hum
- De-click
- De-crackle
- De-clip
- Breath Control
- De-plosive
- De-ess

UTILITY MODULES

- EQ Filters
- Signal Generator
- Phase
- Azimuth
- Center Extract
- Resample
- Stereo Width
- Channels

PROCESS MODULES

- Distance
- Reverb
- Dynamics
- Sound Enhance
- Fade
- Gain
- DC Offset
- Loudness
- Control
- Trim
- Silence
- Harmonic Clarity
- Time & Pitch
- Music Tailor
- Loop Creator
- Dither
- Plugin Host
- Match

AI MODULES

- Stem
- Remix
- Voice Isolate
- Dialogue Rebalance
- AI Speech
- De-noise
- Spectral Recovery
- Voice Enhance
- Transcribe
- Voice Rebuild
- Prompt Isolate
- AI Retouch
- Generate
- Text to Speech

Tools & Workflows

- Batch Processing
- Markers
- Statistics & Analysis
- The Notes Editor (Pitch & MIDI)
- The Text Editor

Exporting

- Saving & Exporting

Reference

- Keyboard Shortcuts
- AI Features & On-Device Models
- Settings, License & Updates

PART

Getting Started

Welcome to FFT

FFT is a native audio editor for macOS and Windows, built for cleaning up, repairing and polishing recordings. It pairs a precise waveform editor with a full-colour spectrogram so you can both hear and [see](#) your audio, then fix problems with a deep suite of restoration and processing tools — de-noise, de-hum, de-click, de-clip, de-ess, spectral repair, EQ, dynamics, loudness and more. Editing is non-destructive in feel: every change is undoable, you can audition before you commit, and the original file is never touched until you save or export.

What it is for


FFT is made for anyone who needs a recording to sound its best:

- **Podcasters and dialogue editors** — strip out hum, hiss, mouth noises, plosives and room reverb, and hit a target loudness.
- **Musicians and producers** — repair clicks and clipping, sculpt tone, manage stereo width, and master against a reference.
- **Post and restoration users** — circle a cough or a chair squeak on the spectrogram and rebuild it from the audio around it.

Everything runs on your computer

FFT does all of its work locally. The audio editing and DSP run on your machine, and the AI tools (transcription, speech de-noise, note detection and the rest) run on-device using your computer's GPU — the Apple Silicon GPU and Neural Engine on macOS, your GPU (or CPU) on Windows — nothing is uploaded to a server to be processed. Some AI models are bundled with the app; larger optional engines must be installed locally before use, then stay on your computer. Where an AI feature is built on an outside model, its panel shows a "Credits & License" link.

A 60-second tour of the main window



When you launch FFT with nothing open, the workspace invites you to **Open or drag an audio file to begin** — click **Open Audio File...**, drag a file onto the window, or press **⌘O** (Ctrl+O on Windows). Once a file is loaded, four areas frame your work:

- **The top toolbar** holds your most-used actions left to right: Open and Save; Cut, Copy, Paste and Delete; Undo and Redo; the quick Silence, Normalize and Reverse processors; and the **selection-tool picker** for choosing what you drag-select (a time range, or — on the spectrogram — a box, frequency band, or freehand lasso region). On the right you'll find the **view-mode picker**, spectrogram settings, statistics, a performance readout, and the button that shows or hides the restoration panel.
- **The main display** is your audio. Use the view-mode picker to switch between **Waveform** (amplitude over time), **Spectrogram** (frequency content as colour), **Split** (waveform above, spectrogram below), **Notes** (a pitch editor) and **Text** (a transcript editor). FFT reveals the spectrogram automatically when a tool or module needs it.
- **The restoration panel** opens along the right edge and lists every processing module — 45 in all, grouped into **Repair** (11, from Spectral Repair and De-click to De-ess), **Utility** (8), **Process** (14, including Sound Enhance, Music Tailor and the unified Match) and **AI** (12, including Voice Isolate, Dialogue Rebalance and Text to Speech). Pick one and it opens as a floating tool window you can drag anywhere, preview, and apply.
- **The bottom transport bar** runs playback and metering: jump to start, Play/Pause (**Space**), Stop, **Loop** the selection, **Record** from an input, and an **A/B** button to flip between the original and your last edit. It also shows the playhead time and selection length, live level meters and momentary loudness, a playback **volume** slider (which never changes the file itself), and zoom controls for both timeline and amplitude.

A few things worth knowing early: edits are fully **undoable** (**⌘Z** / **⇧⌘Z**, or Ctrl+Z / Ctrl+Shift+Z on Windows), the title bar marks a file "Edited" until you save, and the menu bar mirrors the toolbar with extra commands under **Process**, **Generate**, **Restoration**, **Transport** and **Markers & Regions**.

Take the interactive tour

Help ▶ Interactive Showcase... opens a tour of FFT's headline features — Spectral Repair, Spectral Recovery, Generate, Stem Remix, Voice Isolate, Dialogue Rebalance, Text Mode, Match, Music Tailor and the Notes Editor — each narrated over a short demo, with a **Try it in FFT** button that opens the real tool on the demo audio. It is the quickest way to see what a module does before you read its chapter.

Each chapter works the way you would. You hear the problem; the demo audio then stops while the cursor sets the edit up — drawing a selection, picking an engine tab, typing a prompt — and the sound only changes when the cursor presses the button that does the work (**Process**, **Render**, **Rebuild**, **Analyze...**). The audio line reads **AUDIO PAUSED** during those quiet working steps; the A/B chips below play any take you want to hear, whenever you want to hear it.

The tour runs at your pace. The **Pace** switch next to **Play tour** offers two settings. In **Guided** (the default) each step plays its demo and then waits for you: the line under the stage tells you what you are looking at, and a large **Continue ▶** button appears with its own **Space** key cap and a soft pulsing ring — press it, or **Space**, **→** or **Return**, when you are ready for the next step. **←** goes back one step and plays it again. A step that is playing audio keeps looping while it waits, so you can listen as long as

you like, flip between the A/B chips, or move the faders. In **Auto** the tour plays itself straight through, moving on as soon as each step's demo has run, and rolls into the next chapter at the end. Your choice is remembered the next time you open the showcase.

How this manual is organized

After this introduction, the manual walks through FFT in the order you'll meet it:

- **The basics** — opening, navigating, selecting, playing, recording and saving.
- **Editing** — cut/copy/paste, silence, gain, normalize, fades, reverse and channel work.
- **The modules** — a chapter per restoration and processing tool, from De-noise and De-hum to Spectral Repair, EQ, Dynamics, Loudness and the AI tools.
- **Working faster** — module chains, presets, batch processing across many files, and keyboard shortcuts.


Each module chapter explains what the tool does, when to reach for it, every control on its panel, and a short step-by-step. Skim what you need, or read straight through.

macOS and Windows

FFT is the same application on macOS and Windows, and this manual covers both. Keyboard shortcuts are given for both platforms — the macOS chord first, then the Windows one, as in **⌘R** (Ctrl+R on Windows); in the shortcut tables they sit side by side. Where a step names a system feature the two platforms call differently — Finder (macOS) or File Explorer (Windows), System Settings or Windows Settings — both names are given. The screenshots were captured on macOS; on Windows the window frame, menus and file dialogs differ cosmetically, but the same panels, modules and editor views are there. The few places where the platforms genuinely differ say so where they come up: Audio Unit versus VST3 plug-ins, which export formats each platform can write, the languages some voice and lyric features cover, and — on Windows — a few AI models that run on the CPU rather than the GPU because that measured faster or more reliable.

The FFT Window

This chapter is a guided tour of the main window. Once you know where everything lives, the rest of the manual will feel obvious. From top to bottom you have the **document tabs**, the **toolbar**, the editor area with its **rulers** and **overview strip**, the **selection readout bar**, and the **transport bar** with its **level meters**.



Document tabs

Every open file gets its own tab in the strip across the very top. Tabs work just like a browser:

- **Switch files** by clicking a tab. Switching is instant — FFT keeps each open file fully loaded so there is no reload wait.
- **Open a fresh empty tab** with the **+** button on the right.
- **Close a tab** with the small **×** on its left. A faint coloured dot on the right of a tab name means that file has **unsaved changes**.
- **Hover** a tab to see its full file path. The active tab's name is highlighted.

You can also drag an audio file from Finder (macOS) or File Explorer (Windows) straight onto the window: the first file lands in the current empty tab, and each additional file opens in its own tab.

The toolbar


The toolbar runs along the top, just under the tabs. From left to right it groups the actions you reach for most:

- **File** — Open (⌘O, or Ctrl+O on Windows) and Save (⌘S, or Ctrl+S on Windows).
- **Clipboard** — Cut (⌘X, or Ctrl+X on Windows), Copy (⌘C, or Ctrl+C on Windows), Paste (⌘V, or Ctrl+V on Windows), and Delete. Buttons dim when they don't apply (for example, Paste only lights up when there's something to paste).
- **History** — Undo (⌘Z, or Ctrl+Z on Windows) and Redo (⇧⌘Z, or Ctrl+Shift+Z on Windows).
- **Quick process** — Silence, Normalize and Reverse, applied to your selection or the whole file.
- **Selection tool** — a segmented control for choosing how you select: Time, Box (time × frequency), Frequency band, and Lasso (freehand). Hover any segment for a one-line description.
- **Instant Process** — a small pop-up menu that turns one-click processing on or off and picks what it does: **Attenuate**, **De-click**, **Fade**, **Gain** or **Replace**. While it's on, every selection you finish on the spectrogram is processed immediately with the chosen mode — no panel needed. The **I** key toggles it on and off (also at Process ▸ Instant Process).

- **Spectral repair** — Heal and Attenuate buttons that preview a fix on the spectrogram; when a preview is live you also get Bypass/Audition, Apply and Cancel buttons.
- **Text-mode extras** — in the Text view the toolbar gains a **Transcription** button (opens the Transcribe module to re-run the transcription) and a miniature transcript locator (see *The Text Editor*).

On the right side of the toolbar are the **view controls**: the five-way view picker (covered next), the **Spectrogram Settings** menu, the **Statistics** and **Performance** buttons, and a **panel toggle** that shows or hides the module panel on the right edge of the window.

The module panel on the right edge is the **sidebar** — from top to bottom, the **History** panel (your edit timeline), the searchable **module launcher** grouped into Repair, Utility, Process and AI, and the **Module Chain** you're building.




The five view modes

The segmented picker near the right of the toolbar switches how the editor draws your audio. Hover each icon for a tooltip.

- **Waveform** (⌘1, or Ctrl+1 on Windows) — amplitude over time, the classic blue waveform.
- **Spectrogram** (⌘2, or Ctrl+2 on Windows) — frequency content over time, drawn as a colour map.
- **Split** (⌘3, or Ctrl+3 on Windows) — waveform on top, spectrogram below, sharing one timeline (**Waveform + Spectrogram** in the View menu).
- **Notes** (⌘4, or Ctrl+4 on Windows) — a piano-style pitch editor for note-level editing. When you leave the Notes view with unapplied changes, FFT asks whether to **Keep**, **Discard**, or stay and keep editing.
- **Text** (⌘5, or Ctrl+5 on Windows) — a transcript editor: edit the transcribed words and FFT rebuilds the changed audio with a synthesised voice (see *The Text Editor*).

You can also switch from the **View** menu, which mirrors these choices.

View > Show Chords adds a chord band between the overview strip and the editor: FFT analyses the whole file on demand (the same chord detector Music Tailor uses) and names each chord in green along the timeline, re-analysing whenever you edit the audio. Drag the band's lower edge to make the chord names bigger. It is an inspection tool, so it starts hidden in every session.



Rulers and the overview strip

The editor wraps your audio in scales so you always know where you are:

- **Time ruler** — runs along the top *and* bottom of the display, labelled in minutes:seconds (and finer when you zoom in). Click or drag on a time ruler to scrub the playhead. You can switch the readout to other formats, including samples, from the View menu's **Time Format** submenu. Green flags mark any position markers you've added.
- **Waveform overview strip** — a small whole-file waveform just under the top time ruler. A bright outlined box shows the slice you're currently zoomed into. Drag the box to pan, drag its edges to zoom, click outside it to jump there, and **double-click** to fit the whole file. A scroll over it adjusts the overview's vertical scale.
- **Frequency ruler** — appears down the right side whenever the spectrogram is showing, labelled in Hz/kHz. Drag it up or down to pan the frequency range, scroll to zoom into a band, and **double-click** to reset. When channels are shown separately, the scale repeats for each channel lane.
- **Amplitude ruler** — appears beside the waveform, marking the ± 1 / 0 level lines.

Selection and view readout

Just above the transport, a two-row readout bar reports the exact **Selection** and **View** ranges as editable fields — start, end and length in time, plus low, high and span in frequency. Type a value into any field to set it precisely. This row also holds the **Feather** controls for soft-edged spectral edits and a grip you can drag to your file manager to export your selection as a clip.

The transport bar and level meters

The bar along the bottom is your playback hub.



- **Transport buttons** — Go to start, Play/Pause (Space), Stop, a latching Loop toggle, Record from input, and an **A/B** button to audition the original against your last edit.
- **Time readout** — the playhead position and total duration, in your chosen time format.
- **Selection readout** — the start, end and length of the current selection, or the channel count and sample rate when nothing is selected.
- **Level meters** — green-to-red output bars, one per channel, with a moving peak tick (it turns red on a full-scale peak). While playing, a momentary loudness value in LUFS appears alongside.
- **Volume** — a playback slider that goes up to 150%. This only affects monitoring; it never changes the audio file.
- **Zoom** — buttons for zoom in/out, **Fit**, and amplitude zoom in/out. The same actions live in the View menu, where Fit is ⌘0 (Ctrl+0 on Windows) and Zoom to Selection is ⌘E (Ctrl+E on Windows).

Tip: when there's no file open, the editor area shows an empty placeholder — open or drop a file and the full layout above appears.

Opening, Importing & Recording Audio

Everything in FFT starts with getting audio into a document. You can open existing files, drag them in, interpret headerless data, or record live from an input. Each open file lives in its own tab, so you can work on several recordings at once in a single window.

Starting a new document

When you launch FFT with nothing loaded, the editor shows an empty surface with a large **Open Audio File...** button. Use it to bring up the open panel right away.

- **New File** (⌘N, or Ctrl+N on Windows) opens a fresh, empty tab. The first file you open into an empty, untouched tab is adopted there; any additional files you open each get their own tab.
- **New from Clipboard** (⇧⌘N, or Ctrl+Shift+N on Windows) opens a new tab whose audio is the current clipboard contents (whatever you last Cut or Copied), keeping the clipboard's sample rate and channel count. It's greyed out until the clipboard holds audio.
- **Close File** (⌘W, or Ctrl+W on Windows) closes the current tab. If it has unsaved changes you'll be asked whether to **Save**, **Don't Save**, or **Cancel** first.
- **Close All...** (⇧⌘W, or Ctrl+Alt+W on Windows) closes every tab after one confirmation. If any files have unsaved changes the dialog says how many and offers **Save All & Close**, **Close Without Saving**, or **Cancel**.
- **Save All** (⇧⌘S, or Ctrl+Alt+S on Windows) saves every open tab that has unsaved changes through the same path as **Save**.

Opening files

Choose **File ▶ Open...** (⌘O, or Ctrl+O on Windows) and pick a file, or simply **drag an audio file onto the window**. Both routes behave the same way: the first file lands in the current empty tab, and extras open in new tabs.

FFT reads a wide range of formats out of the box, including WAV, BWF, AIFF/AIFF-C, CAF, Sound Designer II, Apple's AU/SND, W64 and RF64, MP3 and MP2, AAC/ADTS, M4A/M4B/M4R and MP4 audio, AC3, AMR, 3GP, FLAC, Apple Lossless, the classic VOC/8SVX/IFF formats, and native Ogg Vorbis, Ogg, and Opus. The release app also bundles **ffmpeg** inside itself (nothing to install), so many further formats work out of the box too — such as WMA, Monkey's Audio, WavPack, Musepack, TTA, Speex, Matroska/WebM audio, DSD, and more. (In a development build without the bundled helper, FFT falls back to an `ffmpeg` on your `PATH`, and installing one unlocks those formats.)

Whatever the source format, the audio is brought in at its own native sample rate and channel layout. Files you open are added to **File ▶ Open Recent** for quick reopening; that menu also offers **Clear Menu** to empty the list.

Tip: Raw imports (described below) are deliberately kept out of Open Recent, because reopening them through the normal decoder would misread the data.

Downloading audio from a URL

Choose **File ▶ Download Audio...** (⌘⌘D, or Ctrl+Shift+D on Windows) to fetch audio straight from the internet. Paste a link into the panel's **URL** field — a direct link to an audio file, a web page that contains audio (FFT scrapes the page for it), or a video/streaming site (handled through the built-in **yt-dlp** downloader that ships inside the app — nothing to install) — and click **Download**. A status line and progress bar track the fetch, and **Cancel** aborts it mid-way. The finished download is stored in FFT's local data folder (`~/Library/Application Support/Fourier` on macOS, `%LOCALAPPDATA%\Fourier` on Windows) and opens as a document, exactly like a dropped file.

Importing raw (headerless) audio

Some files contain sample data with no header describing how to read it. Choose **File ▶ Import Raw Data...**, pick the file, and FFT opens a dialog where you tell it how to interpret the bytes:

- **Encoding** — the sample format (for example 16-bit integer, and the other supported encodings).
- **Byte order** — little- or big-endian.
- **Channels** — how many interleaved channels the data holds.
- **Sample rate (Hz)** — the playback rate to assume.
- **Offset (bytes)** — how many leading bytes to skip before the audio begins.

When you open the file, FFT inspects the first part of it and makes an educated first guess at the encoding and byte order. You can also press **Detect** at any time to re-run that analysis using your current channel and offset values. Click **Import** to load with the chosen settings, or **Cancel** to back out. Because a raw import has no original file format, **Save** behaves as **Save As / Export** afterwards.

Recording from an input

To capture live audio, choose **Transport ▶ Record** (⌘R, or Ctrl+R on Windows) or click the **Record** button in the transport bar (it has a "Record from input" tooltip). The first time you record, the system asks for microphone permission; if it's denied, FFT tells you to enable it in System Settings ▶ Privacy & Security ▶ Microphone (macOS) or Settings ▶ Privacy & security ▶ Microphone (Windows). Recording captures from the input device chosen in **Settings ▶ General ▶ Playback ▶ Input device** (the default, **System Default**, is your computer's current input), at that device's own sample rate and channel count. The choice takes effect the next time you start recording — and if the chosen device has since been unplugged, FFT falls back to the system default and says so.

- If the document is **empty**, the recording becomes the document and the view zooms to fit the new take.
- If the document **already has audio**, the recording is inserted at the playback cursor. If its rate differs from the document, it's matched automatically on the way in.

Press **Record** again (the menu item now reads **Stop Recording**) to finish. If playback was running when you started, it pauses for the take. Should your input device be unplugged or change its rate mid-recording, capture ends gracefully and the audio recorded up to that point is kept rather than lost.

Very long files stream automatically

Opening a very large file (roughly 100 MB or more) switches FFT into a streaming mode: instead of loading the whole recording into memory, it reads exactly the portions it needs from disk. You can browse, play, and work with these long files without the heavy memory cost of decoding them all at once.

Reopening your last session

On a normal launch, FFT reopens the files you had open when you last quit (missing files are simply skipped), so you can pick up where you left off. If the app didn't quit normally last time and you had unsaved changes, you'll first be offered the chance to **Restore** that work or **Discard** it — restored documents come back with their original names, still marked as having unsaved changes.

PART

Playback, Navigation & Editing

Playback & the Transport Bar

The transport bar runs along the bottom of every editor window. It holds the playback controls, a live time readout, the selection readout, the output meter, the playback volume, and the zoom buttons. Nothing here changes your audio file — it only controls what you hear and how you move through the recording.



The transport buttons

Working left to right, the cluster of buttons at the far left controls playback:

- **Go to start** — jumps the play cursor straight back to the beginning of the file.
- **Play / Pause** — starts playback, or pauses it. When you pause, the cursor stays exactly where the sound stopped, so the next play resumes from that point rather than snapping back. The button shows a play icon when stopped and a pause icon while playing, and highlights while audio is running. **Shortcut: Space.**
- **Stop** — halts playback and leaves the cursor where the sound stopped (so you can play again from there). Available only while something is playing. You'll find it as **Transport ▸ Stop** in the menu bar too.
- **Loop** — a toggle that stays on until you switch it off. While it's lit, playback loops: reaching the end of the audio restarts from the beginning, and when you have a time selection, playback loops inside the selection instead. Pausing and resuming (Space) keeps the loop; switching the toggle off mid-play lets the current pass finish and stop at the loop's end. It's also at **Transport ▸ Loop**.
- **Record** — captures audio from your input. The icon turns into a red stop badge while recording; press it again to finish. The first time you record, FFT asks for microphone permission. When you stop, the captured audio is dropped in at the cursor — or, if the document was empty, it becomes the whole file. **Shortcut: ⌘R** (Ctrl+R on Windows) (the menu item reads **Record** or **Stop Recording** depending on state).
- **A / B bypass** — auditions your audio with and without the most recent edit. Press it to flip between "A" (the original, pre-edit sound) and "B" (your edited result) so you can hear exactly what an effect did. The button is only available when there's a recent edit to compare against, and the letter it shows is whichever you're hearing right now — **B** for your edited result, **A** while you're auditioning the original. It's in the menu too, as **Transport ▸ A/B: Hear Original** (or **A/B: Back to Edited** once you've flipped), with the shortcut **Opt-B**.

Tip: Play/Pause respects your selection. If you have a region selected and press Play, FFT plays just that region; with no selection it plays from the cursor onward.

Soloing a single channel

On a multichannel file shown with **View ▶ Show Channels Separately**, single-clicking a channel's lane label solos that channel. While exactly one channel is soloed, the transport plays **only that source channel** — everything else is silent — so you can audition one channel in isolation. Click the soloed label again (or return to the combined view) to hear all channels again; with no channel or more than one selected, playback is normal.

How a soloed channel reaches your speakers is set by **View ▶ Play Single Channel as Dual Mono**, a toggle that is remembered between sessions:

- **On** (the default) — the soloed channel is fanned onto every output, so it is centred and audible on both speakers regardless of which channel it is.
- **Off** — the channel is hard-panned to its own native output (a soloed right channel plays only from the right, and so on).

This affects only monitoring, never the file, and is available once a file has two or more channels. Very large files that stream from disk play all channels normally — soloing does not apply there.

The time readout

Next to the buttons, a large monospaced figure shows the current play-cursor position, followed by a smaller figure after a slash — that's the total length of the file. Both update live as you play.

You choose how this time is displayed from **View ▶ Time Format**. The current choice is marked with a check. Your options are:

- **Samples** — the raw sample count.
- **Time (h^M s)** — minutes, seconds and milliseconds (hours appear once the file is long enough).
- **Timecode (24, 25 or 30 fps)** — shown as hours:minutes:seconds:frames.

The format you pick is remembered between sessions and is also used by the time ruler and the selection readout, so everything stays consistent.

The selection readout

Just past the time, a small line of text describes your current selection: its start, its end, and its length in parentheses, all in your chosen time format. When nothing is selected, this area instead shows the file's channel count and sample rate — a handy at-a-glance reminder of the format you're working in.

Output meter and loudness

On the right sits the **output level meter** — one horizontal bar per channel, filling green through yellow to red as the signal gets louder, with a moving tick that marks the recent peak. The tick turns red when a channel reaches full scale, warning you of potential clipping. Mono files show a single bar; stereo and multichannel files show one bar each (up to eight).

While audio is playing, a small **LUFS** figure appears beside the meter showing momentary loudness — a perceptual loudness reading taken over a short rolling window. It disappears when playback stops.

Volume

A small slider, flanked by speaker icons, sets the **playback volume**. It ranges from silence up to 150 percent, so you can boost quiet material for monitoring. This affects only what you hear — it never alters the audio file itself. Your setting is saved and restored the next time you launch the app.

Zoom

The buttons at the far right control how much of the waveform or spectrogram you see:

- **Zoom out** and **Zoom in** — step the horizontal (time) zoom in and out, keeping the centre of the view in place.
- **Fit** — scales the whole file to fill the window.
- **Amplitude zoom in** and **Amplitude zoom out** — stretch or shrink the waveform vertically, making quiet detail easier to see without changing the audio.

Tip: The view modes themselves are switched from the **View** menu — **⌘1** (Ctrl+1 on Windows) for waveform, **⌘2** (Ctrl+2 on Windows) for spectrogram, **⌘3** (Ctrl+3 on Windows) for the combined view, **⌘4** (Ctrl+4 on Windows) for Notes, **⌘5** (Ctrl+5 on Windows) for Text.

Navigating & Zooming

FFT shows your audio across three linked surfaces — the waveform, the spectrogram, and the time ruler above them. They always move together: zoom or scroll in one and the others follow instantly, so you never lose your place. This chapter covers every way to move around your file and dial in exactly the detail you want.

Horizontal zoom (time)

Horizontal zoom controls how much of the timeline you see at once — from the whole file down to individual samples. All of these live in the **View** menu and work on whichever editor tab is in front.

- **Zoom In** (, or Ctrl+= on Windows) and **Zoom Out** (, or Ctrl+- on Windows) step the magnification in and out, centred on the middle of the visible area.
- **Zoom to Fit** (, or Ctrl+0 on Windows) frames the entire file edge to edge. This also re-arms automatic fitting, so the view keeps the whole file in sight as you resize the window — until you next zoom or scroll by hand.
- **Zoom to Selection** (, or Ctrl+E on Windows) snaps the view to fill the screen with your current time selection. It is enabled only when you have a range selected.

You can zoom right in to sub-sample detail (useful for pencil edits) or all the way out to the complete file. The view never lets you scroll past the start or end.

Scroll-wheel and trackpad gestures

The waveform and spectrogram share one gesture model, so they always feel identical:

- **Two-finger swipe left/right** (or a horizontal mouse wheel) scrolls along the timeline, complete with the usual trackpad momentum coast.
- **Two-finger swipe up/down** zooms in and out, anchored under the pointer so the spot you're looking at stays put. The coast carries through the zoom too — except when you're zoomed all the way in for single-sample editing, where the coast is dropped so the view holds still under your finger.
- **Pinch** (magnify) zooms around the pointer with no inertia, for precise framing.

A classic notched mouse wheel is treated as coarser steps than a trackpad, so physical wheels still feel responsive.

Vertical (amplitude) zoom

Amplitude zoom magnifies the waveform vertically without changing the audio — handy for inspecting quiet passages or low-level noise.

- **Amplitude Zoom In** (, or Ctrl+Alt+= on Windows) and **Amplitude Zoom Out** (, or Ctrl+Alt+- on Windows) scale the waveform's height. It ranges from normal up to 64× and won't shrink below normal.

This affects only the on-screen waveform drawing; your levels are untouched.

The overview strip

A whole-file overview runs across the top of the editor. It always shows the complete file with your selections and playhead marked, plus a bright box showing the slice you're currently viewing in the main editor. Use it for fast, long jumps:

- **Drag the box** to pan the main view anywhere in the file.
- **Drag either edge of the box** to widen or narrow it — narrowing zooms in, widening zooms out.
- **Click outside the box** to centre the main view on that point.
- **Double-click anywhere** to Zoom to Fit.
- **Scroll over the strip** to adjust the overview's own waveform height (its vertical magnification), independent of the main editor's amplitude zoom.

The playhead

The playhead is the vertical line marking the current position; it appears on the waveform, the spectrogram, and the overview strip.

- **Click anywhere** in the waveform or spectrogram to move the edit cursor (and playback start) to that point.
- **Arrow keys** nudge the cursor left and right; hold **Shift** for a step ten times as large. The nudge is zoom-aware — sample-fine when you're zoomed right in, coarser as you zoom out.
- **Home** jumps the cursor to the very start and **End** to the very end. These jumps scroll their target back into view even while playback is paused, so you never land off-screen.
- **Follow Playhead** (View menu) keeps the moving playhead in view during playback. Toggle it on, and choose a behaviour under **Follow Playhead Mode**:
 - **Page** scrolls a full page ahead each time the playhead reaches the edge — steadier, with fewer interruptions.
 - **Continuous** keeps the playhead centred and the audio scrolling smoothly beneath it.

Frequency zoom in the spectrogram

The spectrogram adds a vertical (frequency) axis you can zoom and pan independently of time:

- **Shift + scroll** pans up and down the frequency axis.
- **Shift + ⌘ (or ⌥) + scroll** (Shift + Ctrl or Alt + scroll on Windows) zooms the frequency range in and out, anchored under the pointer. In separate-channels view it zooms within the lane you point at.
- **Reset Spectrogram to Defaults...** (View menu) restores every spectrogram setting — analysis type, FFT size and window, color map, quality, overlays, frequency range and brightness — after a confirmation dialog. Your audio is not touched. The Spectrogram Settings panel also offers a quick control to undo any frequency zoom.

Tip: Zoom to Selection (E , or Ctrl+E on Windows) followed by a frequency zoom is the fastest way to isolate a single problem — a click, a cough, a stray harmonic — before reaching for a repair tool.

Selecting Audio

Almost everything you do in FFT acts on a selection: cut, gain, normalize, repair, denoise. There are two kinds of selection. A **time selection** picks a span of time across the whole signal and is made on the waveform. A **spectral selection** picks a precise patch of *time and frequency* — a rectangle, a freehand shape, or a frequency band — and is made on the spectrogram. The active selection is always shown numerically in the **selection readout bar** below the editor.

Choosing a selection tool

Pick a tool from the segmented control in the toolbar, or from the **View** menu. Each has a one-key shortcut (no modifier needed):

- **Time** (**T**) — drag on the waveform to select a time range.
- **Box** (**R**) — the Time-Frequency Selection Tool: drag a rectangular time × frequency area on the spectrogram.
- **Freq** (**F**) — pick a frequency band that spans the full time range (or your current time selection).
- **Lasso** (**L**) — freehand-draw any shape on the spectrogram.

Switching to any spectral tool reveals the spectrogram automatically. Hover any tool icon for a reminder of what it does.

Making and adjusting a selection

- **Drag** to create a selection. A plain click without dragging just drops the playback cursor — you have to move a few pixels before a selection is committed, so stray clicks never leave a tiny selection behind.
- **Double-click** the waveform or spectrogram to **Select All**. Use **Edit ▸ Select All** (**⌘A** , or Ctrl+A on Windows) and **Edit ▸ Select None** (**⇧⌘A** , or Ctrl+Shift+A on Windows) to do the same from the menu.
- **Resize** a time selection by dragging either edge; the pointer becomes a left-right arrow when you are over an edge. **Move** a whole selection by dragging its interior (the pointer becomes a hand).
- **Shift-drag** (or shift-click) **extends** the existing selection from its far edge.
- **Box and Freq selections** can be moved or resized the same way once drawn.

Adding to and subtracting from a spectral selection

With the **Box**, **Freq**, or **Lasso** tool, you can build up a compound selection:

- **Shift-drag** adds the new region to what you already have.
- **⌘-drag** or **⌘-click** (Alt-drag or Alt-click on Windows) subtracts the new region from the selection.

This lets you, for example, lasso a noisy patch, then **⌘-draw** (Alt-draw on Windows) to carve a clean note back out of it.

Harmonics

For the spectral tools, turn on **Harmonics** in the readout bar to extend your selection to the harmonics of whatever you selected — the 2×, 3×, and higher multiples of the fundamental. The stepper sets how many harmonics to include (2 to 12). This is perfect for grabbing a whined or hummed note along with its overtones in one gesture. Harmonics is unavailable for the plain Time tool.

Feathering (soft edges)

Spectral edits sound far more natural when their borders fade rather than cut hard. The **Feather** control offers three screen-space presets: **Hard** (no visible feather), **Soft** (a small fixed-pixel fade), and **Wide** (a broader fixed-pixel fade). Because the preset is based on what you see on screen, the softened edge stays visually consistent as you zoom or change the selection size.

Editing the lasso shape

After drawing a Lasso, you can fine-tune the outline: drag any vertex to move it, click on an edge to insert a new point, and drag inside the shape to slide the whole lasso around. Hovering a vertex lets you delete it.

Channel-specific selection

On a multi-channel file, each channel has its own lane in the spectrogram with a small channel tag (L, R, C, and so on). Click a tag to work on just that channel; **Shift-click** more tags to add channels; click the **All** selector (or double-click any tag) to return to all channels.

The selection readout bar

The strip beneath the editor shows your selection precisely and lets you type exact values:

- **T Start / T End / T Len** — the time bounds and duration of the selection.
- **F Low / F High / F Rng** — the frequency bounds and span (spectral selections only).
- A second **View** row mirrors these for the part of the file currently on screen, so you can navigate by typing coordinates.

Edit any field and press Return (or click away) to retarget the selection. When a time selection is active, a small grip appears that you can **drag to Finder (macOS) or File Explorer (Windows)** to export just that span as a WAV clip.

Regions: saving named selections

A **region** is a named time range you can recall at any time — handy for marking verses, takes, or problem spots.

1. Make a time selection.
2. Choose **Markers & Regions > Create Region from Selection** (⌘R, or Ctrl+Shift+R on Windows), or open **Markers & Regions > Regions List...** and click **Create from Selection**.

In the Regions List you can rename a region inline, click its target icon to recall its range as a selection, and delete it. **Next Region** and **Previous Region** in the Markers & Regions menu step through them in order, and **Clear All Regions** removes the lot. **Export CSV...** writes every region's name, start, and end (in seconds and frames) to a file for use elsewhere.

Core Editing

This chapter covers the everyday cut-and-paste editing and the quick one-shot processes that reshape your audio directly: gain, normalize, fades, reverse, invert, insert silence, trim, and sample-level pencil editing. These commands change the audio itself (unlike the non-destructive modules), but the most important thing to know up front is this: **every edit is undoable**. Anything you do here can be reversed step by step, so feel free to experiment.

Most of these commands work on the current selection. If nothing is selected, they apply to the whole file (or, for clipboard operations, to the playback cursor position). Edits respect your active channel selection too — if you have only one channel targeted, the process touches only that channel.

Undo and redo

Choose **Edit > Undo** (⌘Z, or Ctrl+Z on Windows) to step back and **Edit > Redo** (⇧⌘Z, or Ctrl+Shift+Z on Windows) to step forward. The menu shows the name of the next action to undo or redo (for example, "Undo Gain 3.0 dB"), so you always know what you're about to reverse. There is no limit you need to worry about in normal use.

Cut, copy, paste, and delete

These live in the **Edit** menu and use the familiar system shortcuts.

- **Cut** (⌘X, or Ctrl+X on Windows) removes the selected range and places it on FFT's internal clipboard. The audio after the cut shifts left to close the gap.
- **Copy** (⌘C, or Ctrl+C on Windows) copies the selection to the clipboard without changing the file.
- **Paste** (⌘V, or Ctrl+V on Windows) drops the clipboard in. If you have a selection, it replaces that selection; otherwise it inserts at the cursor, shifting later audio to make room. If the clipboard came from a file at a different sample rate, FFT converts it to match automatically.
- **Delete** (Delete key) removes the selected range and closes the gap, just like Cut but without touching the clipboard.

Paste Special (in the Edit menu) gives you precise control over how the clipboard lands:

- **Insert** (⌘⇧V, or Ctrl+Alt+V on Windows) — inserts at the cursor and pushes later audio along. Never overwrites.
- **Replace** (⌘⇧⇧V, or Ctrl+Alt+Shift+V on Windows) — overwrites forward from the selection start (or cursor), running past the end if needed. Later audio is not shifted, so marker positions stay put.
- **Mix** (⌘⇧V, or Ctrl+Shift+V on Windows) — sums the clipboard into your audio at the anchor point, layering the two together.
- **Invert and Mix** (⌘^V, or Ctrl+Shift+I on Windows) — subtracts the clipboard instead of adding it, which is handy for finding the difference between two takes.
- **To Selection** (⌘^⇧V, or Ctrl+Shift+T on Windows) — fits the clipboard into the current selection only, cropping or padding it with silence so the selection length (and every marker after it) stays exactly the same. This needs an active selection.

Tip: Select All (⌘A, or Ctrl+A on Windows) and Select None (⇧⌘A, or Ctrl+Shift+A on Windows) make it quick to act on the whole file or clear a selection before pasting at the cursor.

Quick processes (the Process menu)

- **Normalize (peak)** raises the loudest part of the range up to just under full scale, preserving the balance between channels.
- **Normalize Loudness** matches a target perceived loudness instead of peak level. Pick a preset: **-14 LUFS (streaming)**, **-16 LUFS (podcast)**, or **-23 LUFS (broadcast)**.
- **Silence** zeroes the selected range in place without removing any time.
- **Trim to Selection** crops the file down to the selection, discarding everything outside it.
- **Trim Silence** removes silent passages from the file automatically. For threshold and padding control, use the **Trim Silence** module.
- **Fade In / Fade Out** ramp the level up or down across the selection. (For shape and length control, use the Fade module.)
- **Reverse** plays the selection backwards.
- **Invert Polarity** flips the waveform upside down — useful for phase work and cancellation.
- **Remove DC Offset** recenters the waveform on the zero line, removing any constant offset.

Gain is available from the **Gain** module; like the items above, it's applied to the selection (or whole file) as one undoable step. See the **Gain** module for fixed, peak, and loudness adjustment.

To add empty space, use **Generate > Silence (1s)**, which inserts one second of silence at the cursor or selection start and shifts later audio along.

Pencil editing single samples

When you zoom all the way in — far enough that FFT draws each individual sample as a small dot on the waveform — you can edit the audio sample by sample. Click and drag directly on a sample dot to redraw it: the curve follows your pointer, and dragging across several samples redraws every sample in between in one smooth stroke. The view auto-scrolls if you drag past the edge.

This is the tool for surgically erasing a single click or pop. It's available on mono files, or on a multi-channel file when channels are shown separately (so there's one clear channel to write to). Each pencil stroke is committed as a single undoable edit, so one Undo reverts the whole drag.

Undo & Edit History

FFT keeps a complete, ordered record of everything you do to a file, and lets you move freely backward and forward through it. Nothing you do is ever a dead end: you can step back, change your mind, step forward again, or jump straight to any earlier point in the session. This chapter covers the undo/redo commands, the **History** panel, and how non-destructive previews fit into the picture.

Undo and redo

The familiar commands work exactly as you'd expect, and they live in the **Edit** menu:

- **Undo** — reverses your most recent action. Shortcut: **⌘Z** (Ctrl+Z on Windows).
- **Redo** — re-applies the action you just undid. Shortcut: **⇧⌘Z** (Ctrl+Shift+Z on Windows).

FFT names each step after the action that produced it, so the menu reads back what you're about to reverse — for example **Undo Denoise** or **Undo Fade Out** rather than a generic "Undo". When there's nothing left to reverse, the command simply reads **Undo** and is greyed out; the same applies to **Redo** once you've reached the front of the history.

History tracks more than just changes to the audio itself. Adjusting your selection, adding or moving markers and regions, learning a noise profile, and edits made in the Notes editor are all captured as their own steps, so a single **⌘Z** (Ctrl+Z on Windows) reliably backs out whatever you last did, whatever kind of change it was.

One important rule to remember: if you undo a few steps and then make a *new* edit, the steps you had undone are discarded. The new edit becomes the latest point in history, and there is no longer anything to redo past it. This is standard behaviour for an editor, but it's worth keeping in mind before you branch off in a new direction.

The History panel

The **History** panel sits at the top of the right-hand sidebar, above the module launcher. It shows your whole session as a vertical list, oldest at the top and newest at the bottom, so you can see at a glance everything that's happened to the file.



- The first row is always **Original** — the document exactly as it was when you opened or created it. Undoing all the way back lands you here.
- Each row below is labelled with the edit that produced it, in the order you made them.
- A small marker on the left of each row tells you where you are: a **filled dot** marks your current state, a **hollow circle** marks earlier (undoable) states, and a **dotted circle** marks states ahead of you that you can redo to. Future steps are shown slightly dimmed.

Jumping to any point. Click any row to jump the document straight to that state — no need to press **⌘Z** (Ctrl+Z on Windows) repeatedly. Click an earlier row to roll back to it; click a dimmed (future) row to fast-forward. The panel keeps your current position scrolled into view and highlighted as you move, so you never lose your place. Hovering a row shows a tooltip — *Go back to...* for earlier steps, *Redo to...* for future ones.

The panel stays usable even on an empty document. If you undo all the way back to nothing, the list still holds your redoable steps, and clicking one is the way to bring your work back.

Tip: the History panel and the **⌘Z** / **⇧⌘Z** (Ctrl+Z / Ctrl+Shift+Z on Windows) shortcuts drive the exact same history — use whichever is faster in the moment. The shortcuts are quickest for nudging one step at a time; the panel is best when you want to skip several steps at once or compare where a particular change happened.

A practical limit applies: FFT retains a long run of recent steps (well beyond a typical editing session). Once that limit is reached, the very oldest steps drop off the bottom of the undoable range as new ones are added — so the "Original" row tracks the earliest state still held, not necessarily the moment the file was first opened.

Previews and the undo stack

Most of FFT's restoration and effect modules let you **audition a preview** before committing anything. While a module preview is active, you are hearing the processed result, but the underlying audio is untouched and **no history step is created**. You can tweak settings, toggle **Bypass** to compare against the unprocessed signal, and listen as much as you like without cluttering your history.

A history step is added only when you commit the processing — by pressing **Process** (Apply) in the module footer. At that moment the change becomes a real edit, appears as a new row in the History panel, and can be reversed with **⌘Z** (Ctrl+Z on Windows) like any other.


This separation is deliberate and convenient: previews are for experimenting, history is for committed changes. If you preview a module and then close it, undo a different edit, or start an unrelated action, the un-applied preview is simply discarded — it was never written into your audio, so there's nothing to undo.

PART

The Spectrogram & Spectral Editing

The Spectrogram & Spectral Editing

The spectrogram is FFT's signature view. Instead of showing amplitude over time (the waveform), it shows *frequency content* over time: time runs left to right, frequency runs bottom (low) to top (high), and colour shows how loud each frequency is at each moment. A hum reads as a steady horizontal line, a click as a thin vertical streak, sibilance as a bright cloud up high. Once you can see a problem, you can paint it out directly.



Switch displays with the view picker on the right of the toolbar: **Waveform**, **Spectrogram**, **Split** (waveform above, spectrogram below), **Notes**, and **Text**. The four spectral selection tools auto-reveal the spectrogram when you pick them.

Reading and tuning the spectrogram

Open the settings panel with the **slider** button in the toolbar (or View ▸ Spectrogram Settings, , or Ctrl+Alt+, on Windows). Everything here is display-only — it never changes your audio.

- **Presets** — save the current look with **Save Current...**, recall a saved preset, **Delete** one, or jump back to **Factory Default**.
- **Type** — the analysis engine: *Adaptive Mel* (the default; sharp at both low and high frequencies), *Regular STFT*, *Auto-adjustable STFT*, *Multi-resolution*, and *Adaptively Sparse*.
- **FFT Size** — larger sizes give finer frequency detail but blurrier timing; smaller sizes keep transients crisp. Adaptive types show **Auto** and pick the size from your zoom automatically.
- **Window** — shapes each analysis frame to trade frequency leakage against resolution.
- **Enable reassignment** — sharpens the picture by relocating energy to its true time and frequency.
- **Time overlap / Freq overlap** — higher values (up to 16×) smooth the image in time or across frequency, at greater render cost.
- **Color map** — Cyan to orange, Magma, Grayscale, Ice, or Inferno.

- **High-quality rendering** — leave on for final-quality images; turning it off lets the **Reduce quality above [s]** slider drop to a faster preview on long views.
- **Amplitude range (low / high) [dB]** — the quietest and loudest levels mapped to the colour scale. Raise the low value to hide noise; lower the high value to brighten quiet detail.
- **Cache size [MB]** — memory budget for rendered tiles.
- **Scale** — the vertical frequency mapping: Linear, Mel, Bark, Log, or Extended Log.
- **Frequency** — shows the visible range with a **Reset** button to undo any frequency zoom.
- **Piano roll overlay** and a **Waveform overlay** opacity slider that blends the waveform on top (0% hides it). The frequency ruler stays visible beside the spectrogram.

Tip — the colour/amplitude ruler. Beside the frequency ruler is a vertical colour scale. Drag its top handle to move the ceiling, the bottom handle to move the floor, the middle to shift the whole range, scroll to widen or narrow it, and **double-click** to reset. This is the fastest way to make faint detail pop.

Tip — frequency zoom. Hold **Shift** and scroll over the spectrogram to pan the frequency axis; add **⌘** or **⌥** (Ctrl or Alt on Windows) while scrolling to zoom into a band. You can also drag or scroll on the frequency ruler, and **double-click** the ruler to restore the full range. For stereo, View ▸ **Show Channels Separately** stacks each channel in its own lane. When channels are *not* shown separately, the combined image power-averages the per-channel intensities — a balanced view of all channels: identical channels look exactly like a single channel, and a sound present in only one channel shows at proportionally reduced strength instead of dominating the picture.

Selecting a region to repair

Pick a tool from the segmented tool picker in the toolbar (shortcuts in brackets):

- **Time [T]** — an ordinary time-range selection.
- **Box (R)** — the Time-Frequency Selection Tool: drag a rectangle covering a time × frequency area.
- **Freq [F]** — select a frequency band across the whole time range.
- **Lasso [L]** — freehand-draw an irregular region.

Hold **Shift** while drawing to *add* to the current selection, or **⌥** (Alt on Windows) to *carve out* of it — so you can build up a complex shape. The selection bar also offers **Feather** presets (Hard / Soft / Wide) for soft edges, and a **Harmonics** switch with a count that extends the selection to the multiples of a selected tone. Edit ▸ **Select Harmonics (⌘⌘H, or Ctrl+Shift+H on Windows)** does the same from the menu. With the Lasso, arrow keys nudge the shape (Shift for larger steps) and Delete removes a hovered vertex.

The Spectral Repair workflow

With a region selected, open the **Spectral Repair** panel and choose a tab:

- **Attenuate** — pulls the region down. **Strength** sets how much (in dB). Turn on **Use surrounding audio** to reduce the region toward the level of the audio just before and after it (those reference windows appear bracketed on the spectrogram), with a **Surrounding region** control for how much is sampled.

- **Replace** — rebuilds the region from neighboring texture. **Context source** chooses Both, Preceding or Following; **Surrounding region** sets how much donor audio is sampled; **Strength** blends from a full rebuild down to a subtle fill.
- **Boost** — raises the region by a chosen amount of dB.

Resolution [FFT] sets the analysis size for the repair: larger resolves low frequencies better, smaller keeps transients tight.

Auditioning and committing

1. Select your region (the panel header shows *Region ready* once it can apply).
2. Press **Preview** (or the toolbar **bandage** = heal / **speaker** = attenuate buttons). FFT renders the repaired audio so you can hear it.
3. While previewing, toggle **Bypass** (key **B**) to A/B against the original.
4. Press **Apply Spectral Preview** (or **Return**) to commit it, or **Cancel** (or **Esc**) to discard. Every applied repair is undoable.

Gotcha. Replace and Boost need a *bounded* region (Box or Lasso) — a pure frequency band with no time bounds isn't enough. If a button is dimmed, the panel header tells you what kind of selection is still missing.

PART

The Module Reference

How Modules Work

Almost everything FFT does to your audio happens inside a **module** — De-ess, De-hum, Spectral De-noise, Gain, Reverb, the AI tools, and dozens more. They all share the same anatomy and the same safe, non-destructive workflow, so once you learn one panel you know them all. This chapter covers that shared layout and the habits that make modules fast to work with.

Finding and opening a module


Modules live in the **module launcher** on the right edge of the window. Each row shows the module's name, a one-line subtitle of what it does, and an icon.

- **Filter by category.** A menu at the top of the launcher narrows the list to **All**, **Repair**, **Utility**, **Process**, or **AI**. The rows are grouped under those same category headings.
- **Search.** Type in the **Search modules** field to filter by name or by what a module does — searching "vinyl", "hum", "karaoke", "lufs" or "phone" jumps straight to the right tool. Search ignores accents and case. Click the **x** in the field to clear it.
- **Open it.** Click a row to open that module. Its name highlights while the panel is open. Click the row again to close it, or click once more to bring a hidden panel back to the front.

The launcher is greyed out until you have audio loaded.

The floating panel

Each module opens as its own **floating panel** that hovers above the editor. Because every module owns its panel, you can keep several open at once. Drag a panel anywhere on screen — even off the main window — by its **title bar**; it remembers where you left it. The **x** in the title bar closes it. Panels float only while FFT is the active app and tuck away when you switch apps.



From top to bottom, every panel has the same parts:

1. **Title bar** — the module's name, a spinner (with a **Rendering... n%** read-out on slower renders) while it is working, and the close button.
2. **Preset menu** (left) and the **?** help badge (right).
3. **The controls** for that module, plus a per-module **visualization** where relevant (a frequency response curve, a spectrogram, meters, and so on).
4. A live **feedback line** that flags a failed render, or tells you when a render changed nothing.
5. The **footer** of action buttons.

Presets and the "?" help badge

- **Presets.** Open the **Presets** menu to pick a built-in **Factory** preset or one of **My Presets**. When you have changed the controls away from the loaded preset, an asterisk (*****) appears next to its name. Click the **+** beside the menu to save the current settings as a new named preset. With one of your own presets selected, the menu also offers **Update** (overwrite it with the current settings) and **Delete**.
- **Help.** Hover the **?** badge for a quick tooltip, or click it for a short popover describing exactly what the module does.

The footer: Preview, Bypass, Process

The footer is where you audition and commit. Left to right:

- **Preview** (Shift+Space) — renders a non-destructive audition and plays it, so you hear the module's effect without touching the file. Preview is a **listen toggle**: the first press renders and starts playing; press it again while the audition is sounding and it **stops the audio** but keeps the rendered preview; press once more to re-listen without re-rendering. Pressed while a render is still in flight, it cancels that render. The button lights up while a preview is live.
- **Bypass** (Shift+B) — while a preview is auditioning, toggles between the processed and the original signal for a quick A/B. The button stays lit while you are hearing the unprocessed audio, and is greyed out until there is a live preview to compare against.
- **The + overflow** — a small menu offering **Cancel Preview** (discard a preview that is rendering or auditioning) and, for supported modules, **Add to Module Chain** (stack this module's current settings into a saved processing chain — see the Module Chain chapter). The button appears only when the module can join a chain or while a preview is live; otherwise it is hidden.
- **Reset to defaults** — the circular-arrow button returns every control to the module's default settings. It is disabled once you are already at the defaults, and hidden entirely on the few modules that have no default snapshot (such as Channels).
- **Update step** — appears only while a panel is editing a saved Module Chain step; it writes your live tweaks back into that step (see the Module Chain chapter).
- **Process** — applies the processing for real, writing it into the file as a single undoable edit. Its label stays a plain **Process** whether or not a preview is live; when a matching preview has already been rendered, Process simply commits that render. One module puts the commit on its own in-panel buttons — Phase, which shows no Process pill, commits each of its polarity, delay and rotation operations from its own Apply button.

Auditioning and committing, and what gets processed

Dial the module in, press **Preview** (Shift+Space) to hear it, and — when it sounds right — press **Process** to commit. While the preview plays, **Bypass** (Shift+B) A/Bs it against the untouched signal, and pressing Preview again stops the audio without discarding the render. Slower renders show a **Rendering... n%** read-out in the title bar; some modules stream the preview from the playhead, so you can start listening while the rest is still rendering. Because every commit lands as a single undoable edit, the working rhythm is *preview, compare, commit, step back if needed* — the **A/B button in the transport bar** re-auditions the audio as it was before the last edit, and **Undo** (⌘Z, or Ctrl+Z on Windows) removes the render entirely. The feedback line above the footer flags a failed render, and tells you when a render changed nothing — that's a clean "nothing to repair", not a failure.


Spectral Repair keeps its own extra entrypoints — its previews can also be started from the spectrogram toolbar, the Restoration menu's Preview commands, or Instant Process (see the spectral editing chapter) — but the footer's **Preview** (Shift+Space), **Bypass** (Shift+B) and **Process** now drive it just like every other module.

One rule governs the scope of processing across modules: **if you have an active selection, the module processes only that selection; with nothing selected, it processes the whole file.** So select a noisy passage to clean just that region, or click away to clear the selection and treat the entire take. (A few modules are inherently selection-based — for example, gap-fillers need you to select the gap first.)

A handful of analysis and generation tools (such as Transcribe and Prompt Isolate) drive their own workflow with dedicated buttons and don't show the standard footer at all. Their panels explain their own steps.

The Module Chain

Most of the time you reach for a single module: open it, dial it in, hit **Process**, done. But cleanup work is rarely one step. A noisy interview might want de-hum, then de-noise, then a touch of EQ, then a loudness pass — every time, in the same order. The **Module Chain** is where you assemble that sequence once and run it as a single, undoable pass.



A chain is an ordered list of steps. When you run it, FFT processes them top to bottom, feeding the output of each step into the next, then commits the whole result as one entry in your undo history. The chain runs over your current selection (or the whole file if nothing is selected) and across the channels you have selected, exactly like an individual module.

Building a chain

Each module captures its own settings as a step, so you tune a module first and then add it:

- **Add a module to the chain.** With a module open and configured the way you want, click the small **+** button in the module's footer and choose **Add to Module Chain**. The step is appended to the bottom of the chain with the exact settings you have dialed in at that moment.
- **Or add straight from the launcher.** Every chainable module's row in the module launcher carries a small stack-plus icon on its right; clicking it appends that module to the chain without opening its panel — with default settings (Gain and Fade capture their current settings), ready to be tuned in place (see editing steps below).

- **Gain and Fade have a shortcut.** Both of those modules show a dedicated **Add to Chain** button next to their Apply button — same result, one click instead of two.
- **Snapshot, not a live link.** A step remembers the settings it was added with. If you later change the module's controls, the step you already added does not change — edit the step itself instead (below).
- **Learned profiles travel with the step.** Modules that learn from your audio (for example a noise profile or an EQ match reference) bake what they learned into the step at the moment you add it. That step keeps using its own captured reference even if you re-learn or clear the live one later.

Reordering, bypassing, and editing steps

The chain list shows one row per step with its name and a short summary of its settings. Each row has controls on it:

- **Move up / Move down.** Use the up- and down-chevron buttons to change a step's position. Order matters — de-clicking before de-noising is not the same as the reverse.
- **Bypass a step.** Click the power button to toggle a step off without deleting it. A bypassed step dims and is skipped when the chain runs, so you can A/B whether it is earning its place. Click the power button again to re-enable it.
- **Remove a step.** Click the ✕ on the row to delete just that step.
- **Edit a step's settings.** Click the sliders icon at the left of the row to open that step in the module's real floating panel, seeded with the step's stored settings. Tweak the controls live, then click **Update step** in the panel's footer to save the changes back into the stored step — closing the panel without updating leaves the step as it was.

Running the chain

Click **Apply Chain** to render every enabled step in sequence and commit the result as a single undoable edit. A progress panel reports which step is running ("Step 2 / 4 ...") and offers a working **Cancel** — cancelling aborts cleanly and touches nothing. If a step finds nothing to change, the chain simply moves on; if some steps could not run, the undo entry says so (for example "Module Chain (3 of 4 steps)"). After a run, a **Last run** report under the steps lists each one as Applied, No change, Skipped, or Failed, so a quiet skip never goes unnoticed.

- **Clear chain.** The trash button empties the whole chain.
- **Saved chains and presets.** The stacked-squares menu is the chain preset hub. **Save Chain...** stores the current chain under a name of your choosing; your saved chains then appear in a **My Chains** section, where each can be loaded, renamed, or deleted. A **Built-in** section offers three read-only starting chains — **Podcast Fade + LUFS**, **Broadcast Peak Fade**, and **Soft Head/Tail** — which you can load and then reorder or trim. Loading any chain, built-in or saved, replaces the current one.

When to chain, and what can't be chained

Reach for a chain when you have a repeatable, multi-step workflow — especially one you also want to run across many files. In the **Batch** window, a **Module Chain** toggle runs the batch-runnable parts of your saved chain on every file after the batch's own stages. Steps that depend on the live editor state, missing learned profiles, unavailable plug-ins, or installed AI engines are marked and skipped instead of being applied blindly.

Most restoration and effect modules can be chained. A specific set cannot, because they either change the shape of the audio in ways a fixed sequence can't safely carry, or they aren't a render-in-place effect at all. The modules that **cannot** be added to a chain are:

- **Spectral Repair** — works on the current spectral selection.
- **Stereo Width** — can change the channel count.
- **Channels** — can change the channel count.
- **Time & Pitch** — changes the audio duration.
- **Resample** — changes the sample rate.
- **Match, in Full Master mode only** — the EQ Curve mode chains fine (see below); the full-master mode, which matches level, width and ceiling against a reference master, does not, and the **+** menu says so until you switch modes.
- **Signal Generator**, **Music Tailor**, and the analysis- and generation-style AI modules — **Transcribe**, **Stem Remix**, **Voice Isolate**, **Dialogue Rebalance**, **Prompt Isolate**, **Voice Rebuild**, **Text to Speech**, **AI Retouch** and **Generate** — have no chain option at all.

For the modules that have a footer, the **+** menu shows a short greyed-out note explaining why instead of an Add option. Run all of these on their own.

The unified Match module: Match chains only in its **EQ Curve** mode, where the step carries the learned reference curve and re-analyzes the *working* audio at run time — so the match reflects what earlier steps have already done, not the pre-chain file. A step saved in **Full Master** mode has no runnable chain equivalent and is dropped rather than half-applied.


Tip: If the chain contains a length-changing step while you have a time selection active, later steps that depend on that selection are reported as skipped rather than applied to the wrong region. Run length-changing work on the whole file, or as its own pass.

Repair Modules

Fix problems in a recording — noise, hum, clicks, clipping, reverb, sibilance and gaps. Most repair modules work on a selection or the whole file and preview non-destructively.

Spectral Repair

Paint over a problem on the spectrogram and rebuild it from the audio around it.



What it does. Spectral Repair fixes a region you mark on the spectrogram by rebuilding it from the surrounding audio. You draw a box, lasso or brush over the offending mark, choose how the gap should be filled, and the module reconstructs that patch of time and frequency from its neighbours so the blemish disappears without punching a hole in the sound. It works in five modes — Auto (let the module choose), Attenuate (pull down), Replace (rebuild from donor texture), Regenerate (re-render with the on-device generative engine) and Boost (lift). Whichever mode reads surrounding audio shows exactly the windows it samples as brackets on the spectrogram, and those brackets follow your selection live while you drag it.

When to use it. Reach for it to erase isolated noises that you can see and circle on the spectrogram: a cough or chair squeak over speech, a stray bird chirp, a phone-ring artefact, a digital glitch, or a cymbal bleed sitting on top of a vocal note. Use **Regenerate** when the interpolating repairs smear — dropouts, dense click bursts, crushed clipping — and you want the span re-imagined from its context instead of patched bin-by-bin.

Controls

- **Mode tabs (Auto / Attenuate / Replace / Regenerate / Boost)** — pick the kind of repair. Each tab reveals its own controls below.
- **Auto** (tab) — hands the choice to the module. It renders every classical candidate for the marked region on a scratch copy — Attenuate at two depths, Replace with each fill — measures how well each result matches what the surrounding audio predicts, and keeps the winner. There is nothing to set: the point of the tab is that it chooses the tool *and* its settings, per channel, for this one region. That costs about half a second per channel, which is why it is a tab you pick rather than a silent default for the others. The readout names the tool it landed on, Preview auditions exactly that result, and Render commits it. Auto needs a single box region — lasso and compound selections use the other tabs.
- **Strength** (Attenuate) — how hard the marked region is pulled down, from 0 (gentle) to 100 (strongest); the default is 37. When *Use surrounding audio* is on, this sets the *maximum* reduction rather than a flat cut — with one deliberate exception: where the module can see that the defect in a given spot is louder than that limit allows it to remove, it goes further, but never below the level of the clean audio around it. So the marked region is pulled down to its neighbourhood rather than left with an audible remnant just because the dial was set low.
- **Use surrounding audio** (Attenuate) — instead of a flat cut, samples the clean audio just before and after the selection and pulls the region down toward that level. The sampled windows are bracketed on the spectrogram and reshape live as you adjust the length or move the selection.
- **Direction** (Attenuate) — how the surrounding audio guides the attenuation: **2D** (both axes), **Horizontal** (along time) or **Vertical** (along frequency). Always visible; it disables when *Use surrounding audio* is off, because a flat cut has no direction.
- **Context source** (Attenuate, Replace and Regenerate) — choose **Both**, **Preceding** or **Following**. Only the selected flank(s) condition the repair, and only those brackets appear. A source with no audio at the file edge is rejected instead of silently running without useful context.
- **Surrounding region** (Attenuate, Replace and Regenerate) — how much audio on each chosen side is sampled (0.10–5.00 s). Attenuate and Replace default to 0.50 s; Regenerate defaults to 5.00 s. The spectrogram brackets resize continuously while you drag the slider.
- **Strength** (Replace) — 1.0 fully replaces the region with the synthesized surrounding texture; lower values blend more of the original through (0.05–1.00 mix, default 1.0). Replace samples robustly across the chosen surrounding window so one unrelated event farther along that flank does not dominate the fill.
- **Fill** (Replace) — how the region is rebuilt. **Similar audio** (default) hunts for a closely matching passage in the surrounding audio and grafts it over the selection, falling back to a smooth rebuild when nothing similar is found. **Smooth** always rebuilds from the surrounding spectral envelope only — no grafted passage and no predictive continuation — and writes strictly inside the frequency band you marked, so it leaves the rest of the spectrum untouched. Reach for **Smooth** on noise-like or dense material where a grafted copy would repeat audibly.

- **Regenerate** (tab) — Click **Process** to re-render the selected span with the on-device generative engine. The source and duration controls determine the full-band flanking audio used to condition the result; a spectral lasso still confines the generated change to the drawn shape. The quality recipe remains pinned to the validated settings (Stable Audio 3, full replacement strength, eight steps, no text prompt) so the defect is not leaked back into the fill. There is no preview: the result applies directly as an undoable edit — run it again for a different take. Selections under the engine's minimum span (~0.19 s) are widened to it, centered on the selection; the caption discloses this. Requires the Stable Audio 3 engine (install it from the Generate module).
- **Boost** (Boost) — raises the selected region (0 to 18 dB, default 6).
- **Heal direction** — the interpolation direction used by the separate **Edit ▶ Heal** command and its preview (2D, along time, or along frequency). It does not affect the tabs above.
- **Resolution [FFT]** — **AUTO** chooses the ideal transform at processing time from the painted shape: the more horizontal the mark (long and narrow — a tonal line), the larger the window, up to 8192 for frequency detail; the more vertical (short and broad — a click), the smaller, down to 64 for click-tight timing. The selection's bandwidth also sets a floor so AUTO never picks bins wider than a narrow mark. **AUTO (MULTI-RES)** goes further and searches for this one repair: it scores a ladder of FFT sizes from 256 to 4096 against the surrounding audio, renders the winner, and blends the top two only when they are too close to call — so transients keep their time precision and sustained tones keep their frequency precision. It is slower than a fixed size, and it applies to Replace and to Attenuate with *Use surrounding audio* on a box selection; lasso and compound selections, the Auto tab and the other operations use the standard resolution (2048). Manual sizes pin the resolution instead, from 64 through 8192. This control is hidden on Regenerate because the generative engine does not use an FFT repair size.

How to use it

1. Open the spectrogram and select the problem with the box, lasso or brush tool. The status line confirms when a region is ready.
2. Choose a mode tab and set its controls. The bracketed windows show exactly which surrounding audio the repair will read.
3. Click **Process** to commit the repair; the edit is undoable, so audition the result and re-run with different settings if needed. Note that Spectral Repair cannot be added to the Module Chain — it works on the current spectral selection, so run it on its own.


Tips

- Not sure which mode fits? Start on **Auto** — it renders the classical repairs on your region and keeps the one that measures best, then names the tool it chose, so you can switch to that tab and fine-tune from there.
- Use **Attenuate** with *Use surrounding audio* for marks sitting on busy program: it ducks toward the room rather than cutting a silent notch.
- Drop **Strength** below 1.0 in Replace when a full fill sounds too clean — it lets a little of the original texture through.
- For a defect immediately after a clean phrase, try **Preceding**; for one immediately before a clean phrase, try **Following**. **Both** is the safest default for stationary material.

- Try **Regenerate** on defects that Replace smears — dropouts, dense click bursts, hard clipping. Each run is a fresh take; undo and run again if the first fill isn't right.
- Start with **AUTO** Resolution — it re-reads the painted shape every time you process, covering the full 64–8192 range on its own. Pick a manual size only to pin a specific trade-off. At the small sizes each frequency bin is wide (750 Hz at FFT 64, 48 kHz) — a lasso narrower than one bin cannot be analyzed; the status line says so, and AUTO raises its choice when the bandwidth requires it.

Interpolate

Redraw a short damaged passage by rebuilding the waveform straight across it.



What it does. Interpolate replaces the selected stretch of audio with a fresh waveform that it reconstructs from the clean audio on either side. Rather than simply drawing a straight line through the gap, it studies the slope and rhythm of the surrounding sound and continues that motion inward from both edges, blending the two halves smoothly in the middle. Because it follows the surrounding waveform, it can carry tones, partials and hum through the gap rather than leaving an obvious flat patch.

When to use it. Reach for Interpolate on brief dropouts, digital glitches, mute holes and editing seams — anything from a single sample up to the 4096-sample cap (about 93 ms at 44.1 kHz, 85 ms at 48 kHz). It works best when there is healthy, uncorrupted audio on at least one side of the selection to use as a reference.

Controls

- **Status chip** — a live readout at the top of the panel tells you whether the module is ready. It reads *No audio* (open a file first), *No time selection* (make a selection), *No anchors* (your selection has no clean audio beside it to build from), or *Ready*, and shows the gap length in frames and how many channels will be repaired.
- **Quality** — sets how closely the rebuilt waveform follows the surrounding sound (range 1–8, default 2). At 1 it bridges the gap simply; raising it rebuilds the passage in more detail from the surrounding slopes, which helps on sustained, tonal or harmonically rich material. Very high values on noisy or featureless context can do more harm than good.

How to use it


1. Select just the damaged passage on the waveform — keep the selection tight so untouched audio stays untouched.
2. Confirm the status chip reads **Ready**.
3. Set **Quality** (start at the default of 2).
4. Click **Interpolate** to redraw the selection in place. The edit is undoable, so audition the result and re-run with a different Quality if needed.

Tips

- Keep the selection as short as the damage; the longer the gap, the less convincing the rebuild.
- Selections longer than the 4096-sample limit are refused: clicking **Interpolate** shows a "Selection Too Long to Interpolate" notice with the limit at your sample rate, and the audio is left untouched. Because the cap is fixed in samples, the maximum repairable duration shrinks as the sample rate rises (≈ 46 ms at 88.2 kHz).
- If you see *No anchors*, your selection reaches the very start or end of the file — leave clean audio on at least one side.
- For longer holes or tonal gaps where this leaves an audible scar, try Spectral Repair instead.
- Try Quality 3–5 on held notes or hum, and lower values on percussive or noisy material.

Spectral De-noise

Strip out steady background hiss, hum and air without hollowing out the wanted sound.



What it does. Spectral De-noise reduces broadband noise — tape hiss, air-conditioner rumble, fan noise, electrical buzz, the general "floor" under a recording. It works from a noise profile: a fingerprint of what the noise sounds like across the spectrum. You can teach it that profile from a quiet, noise-only patch with **Learn**, or leave it in adaptive mode, where it estimates the noise on its own. Anything that rises far enough above that profile is treated as wanted signal and kept; the rest is pulled down.

When to use it. Reach for it on dialogue, voiceover, music or field recordings that carry a constant hiss or hum but where the noise stays roughly the same throughout. It is the workhorse for general clean-up before more surgical repairs.

Controls

- **Learn** — captures a noise profile from your current selection. Select a stretch of *noise only* (no speech or music) first, then click Learn. The status line then reads "Profile learned"; with no profile it stays in "Adaptive mode (no profile)" and works from its own estimate. Re-learn any time you find a cleaner noise-only patch.

- **Spectrum display** — a live graph (frequency along the bottom, 100 Hz / 1k / 10k marked). Grey is your selection's spectrum, orange is the learned noise profile, and the coloured line is the profile lifted by the Threshold. Tones poking above that coloured line are kept; everything under it is reduced.
- **Threshold** — how far above the noise profile a tone must rise to count as signal (0–24, default 6). Higher protects more of the wanted sound but lets more noise through; lower digs deeper but risks dulling the source.
- **Reduction** — how hard the detected noise is attenuated, in dB (0–30, default 12). Higher removes more noise; very high can sound artificial.
- **Quality** — stepped A–D, Fast to Best (default B). Higher steps add frequency and time smoothing for fewer artifacts at the cost of speed.
- **Floor** — the residual noise floor (0–0.50, default 0.05). 0 gates the noise away completely; higher leaves a touch of natural ambience so the result breathes instead of sounding gated.

How to use it

1. Select a short, noise-only region and click **Learn** (or skip it to stay adaptive).
2. Select the audio you want cleaned, or leave nothing selected to treat the whole file.
3. Set **Reduction**, **Threshold**, **Quality** and **Floor** while watching the spectrum.
4. **Preview** to audition — if the result sounds muffled or over-thinned, ease Threshold or Reduction back.
5. **Process** to commit, or **Add to Module Chain** to stack it with other modules.

Tips

- Learn from the quietest gap you can find; a cleaner profile means cleaner reduction.
- Push Reduction too far and you get watery "musical noise" — back off and raise Floor slightly to mask it.
- For deep cuts, a higher Quality step (C or D) trades speed for noticeably fewer artifacts.

De-reverb

Attenuates late reverb while preserving the direct signal.



What it does. De-reverb pulls back the room — the ringing, smeared tail that builds up after every sound — while leaving the direct, up-front signal intact. It works by tracking how energy decays in each part of the spectrum: sharp onsets, where the new sound is much louder than the lingering tail, pass through untouched, while sustained, decaying reverberation is turned down. The result is a tighter, drier, more present recording without the obvious gating or hollow tone of cruder tools.

When to use it. Reach for it on dialogue or voice-over recorded in a live room, podcast interviews captured in an untreated space, or any instrument that sounds distant and washed-out. It is a repair tool, not a creative effect — use it to tame excess room, not to make a recording sound dead.

Controls

- **Learn** — Analyses your current selection and estimates how persistent the reverb tail is, then sets **Tail length** automatically and suggests a matching **Reduction** for you. Always a good starting point. If the selection is too short to measure, the panel tells you so; just select a longer passage and try again.
- **Voice mode** — Swaps the engine to a linear-prediction (WPE) dereverb that is stronger on voice and dialogue at moderate reverb, at the cost of much heavier CPU — not recommended for music. While it is on, **Learn** and **Tail length** are disabled (they belong to the standard engine) and **Reduction** acts as the wet/dry blend between the dereverbed signal and the original. Off by default.
- **Reduction** — How aggressively the late reverb is attenuated. Higher values strip away more tail; lower values are gentler and more natural (0–100 %, default 33 %). In Voice mode this becomes the wet/dry blend.
- **Tail length** — How long-lasting the reverb is modelled to be. Higher settings treat the room as more reverberant and reach further into each decay; lower settings target only short, quick tails. Learn sets this from the audio (0.10–0.95, default 0.60). Disabled in Voice mode.

How to use it

1. Select the region you want to clean (or work on the whole file).
2. On voice or dialogue, consider switching **Voice mode** on; otherwise click **Learn** to estimate the tail and seed the controls.
3. Fine-tune **Reduction** and **Tail length** by ear (in Voice mode, just **Reduction** — it is the wet/dry blend).
4. Click **Preview** to audition a non-destructive result.
5. Click **Process** to apply, or use **Add to Module Chain** to stack it with other repairs.

Tips

- If the preview starts sounding dry or thin, or transients lose their edge, lower **Reduction** until only the tail is being pulled back.
- Modest settings sound the most natural — over-reduction can leave a thin, processed quality.
- Run **Learn** on a representative passage (steady speech or a sustained note) for the most accurate tail estimate.
- **Voice mode** shines on roomy dialogue and interviews; on music or dense mixes stick with the standard engine, which is cheaper and safer on non-speech material.

De-hum

Silence mains buzz and its harmonics without smearing the rest of your audio.



What it does. De-hum removes the steady electrical hum that creeps in from mains power — the 50 or 60 Hz tone you hear on recordings made near lights, transformers, or poorly grounded gear. It places a series of zero-phase notch filters at the hum's base frequency and its harmonics, carving out the buzz while leaving the surrounding sound untouched. Because the notches are zero-phase, they remove the tone cleanly without shifting or coloring the audio around them.

When to use it. Reach for De-hum whenever a recording has a constant low buzz or whine — dialogue tracked under fluorescent lighting, a guitar DI with a ground loop, or any take with a steady mains tone running underneath it.

Controls

- **Learn** — analyzes the current selection, detects the hum's fundamental and how many harmonics it carries, and sets Base and Harmonics for you. If no convincing hum is found, it shows "No hum detected."
- **Link** — chooses which harmonics get notched: **All**, **Odd**, or **Even**. The fundamental is always treated. Use Odd or Even when the hum's overtones follow that pattern.

- **Response plot** — shows the input spectrum (gray) overlaid with the notch response (blue dips) on a log-frequency axis, so you can see each notch land on a hum peak.
- **Base** — the fundamental hum frequency to remove (20–500 Hz, default 60 Hz). Two quick buttons set it instantly to **50** or **60** Hz mains.
- **Harmonics** — how many multiples of the base frequency to notch (1–16, default 8). More harmonics catch a buzzier, brighter hum.
- **Filter Q** — notch width (5–100, default 40). Higher Q makes each notch narrower, removing less of the surrounding audio.
- **Slope** — eases notch depth on higher harmonics (0–12 dB/harm, default 0). Higher values treat upper harmonics more gently.
- **Filter DC offset** — also removes a constant DC bias that pushes the waveform off center (off by default).

How to use it


1. Select a region containing clear hum (or the whole file).
2. Click **Learn** to auto-set Base and Harmonics, or set them by hand using the 50/60 buttons.
3. Adjust **Harmonics**, **Filter Q**, and **Slope** while watching the response plot line up with the hum peaks.
4. Use **Preview** to listen — if the notches thin out wanted low end, raise the Q.
5. Click **Process**, or **Add to Module Chain** to stack it with other repairs.

Tips

- Always confirm 50 vs 60 Hz first — picking the wrong mains base leaves the buzz untouched.
- If Process thins out the body of the sound, raise the Q for tighter notches or trim Harmonics.

De-click

Finds and erases the brief pops, ticks and clicks that punch through a recording.



What it does. De-click scans your audio for short impulsive events — vinyl pops, digital glitches, mouth clicks, microphone ticks — and repairs each one by smoothly rebuilding the few samples it spans, so the surrounding tone is left untouched. It only ever touches the sharp outliers it is confident are clicks, never sustained musical content. You can choose a detection method tuned to the kind of clicks you have, and audition exactly what it is catching before you commit.

When to use it. Reach for De-click on digitized vinyl and tape, on spoken-word and dialogue with distracting mouth and lip ticks, and on any recording where isolated pops or glitches break an otherwise clean signal.

Controls

- **Algorithm** — picks the detection method. *Single band* is the straightforward general-purpose detector; *Multi-band (periodic)* runs extra passes tuned for regular, repeating ticks; *Multi-band (random)* merges wider damage and suits scattered vinyl crackle and thumps. (Default: Single band.)
- **Frequency skew** — biases detection toward low or high frequencies. Drag toward **LF** to favour low-frequency vinyl clicks, toward **HF** to favour high-frequency mouth and digital ticks; centred detects across the whole signal. (Range -1.0 to +1.0, default 0.) Active only for the two Multi-band algorithms; it dims for Single band.
- **Sensitivity** — how readily an event counts as a click. Higher catches more and fainter clicks but risks clipping legitimate transients; lower is more conservative. (Range 1.0 to 10.0, default 3.5.)
- **Click widening** — extends the repaired region a little on each side of every detected click, useful for clicks that have a short decay tail rather than a clean impulse. (Range 0 to 3 ms, default 0.)

How to use it

1. Select the region to clean (or leave nothing selected to process the whole file).
2. Pick the **Algorithm** that matches your click type and set **Sensitivity**.
3. Click **Preview** to hear the result without committing — confirm you are removing pops, not music.
4. Adjust, then **Process** to apply. You can also **Add to Module Chain** to stack it with other repairs.

Tips


- If music or speech starts to sound dull or smeared, you are too aggressive — lower Sensitivity or switch to Single band.
- Reserve Click widening for clicks with audible ringing; on clean impulses keep it at 0 to repair the least amount.
- For dense, continuous surface crackle rather than isolated pops, follow up with De-crackle, which mops up the fine residual haze De-click leaves behind.
- Frequency skew only bites on the Multi-band algorithms; select one of those first if you want to steer detection by pitch.

Presets

- **Light Clicks** — gentle single-band cleanup for occasional pops.
- **Vinyl Crackle** — multi-band random detection skewed for record surface noise.
- **Mouth Clicks** — multi-band periodic detection tuned for dialogue lip and mouth ticks.
- **Aggressive** — high-sensitivity multi-band pass for heavily clicked material.

De-crackle

Clears the fine surface crackle and micro-clicks baked into vinyl and optical-film transfers.



What it does. De-crackle targets the dense bed of tiny, single-sample spikes that sits underneath a recording — the constant "frying" texture you hear on vinyl, shellac and optical film, rather than a few loud, discrete pops. It scans the audio, finds samples that jump sharply away from their immediate neighbours, and quietly repairs each one so the underlying waveform is preserved. It works through the material in two passes, so even thick crackle is thinned down progressively without smearing the music underneath.

When to use it. Reach for De-crackle on digitised LPs, 78s and film soundtracks where the noise is a continuous crackly haze across the whole track. Use it after De-click (which handles the bigger, isolated pops) so De-crackle can mop up the fine residual texture that remains.

Controls

- **Threshold** — sets how readily a sample is treated as crackle and repaired. Lower it until the crackle bed audibly disappears; raise it if the audio starts to sound dull or loses fine detail, which means it is also catching wanted transients. (0.5–8, default 2.)

How to use it


1. Select the noisy range, or leave nothing selected to process the whole file.
2. Set **Threshold**, starting around the default and lowering it until the crackle bed vanishes.
3. Click **Preview** to audition the result in place — if it starts sounding dull, raise Threshold.
4. Click **Process** to commit, or **Add to Module Chain** to stack it with other repairs and apply them together.

Tips

- Crackle that survives even low Threshold is usually made of larger clicks — run De-click first, then De-crackle.
- Lower Threshold only as far as the cleaned result keeps its fine detail: push past that and it starts dulling the program.
- For badly worn transfers, render once and run a second light pass rather than forcing one extreme setting.

De-clip

Rebuild the flattened tops of clipped peaks so loud transients breathe again.



What it does. When audio is recorded or rendered too hot, the loudest peaks hit a ceiling and flatten into hard, distorted plateaus. De-clip finds every sample that sits at or above the clip level you set and reconstructs the missing peak above it, drawing a smooth curve through the flattened run instead of the squared-off top. The result restores the natural shape of transients and removes the gritty edge that clipping adds.

When to use it. Reach for it on overdriven location recordings, peaked vocal takes, distorted bass hits, or any file whose waveform shows obviously flat-topped peaks. It works best when only a modest fraction of the signal is clipped; severely mangled material can only be partially recovered.

Controls

- **Threshold histogram + Threshold slider** — the mirrored amplitude plot (dB axis, 0 dB at the outer edges) shows how loud the selected audio is, with two linked lines marking the clip level. Drag the vertical slider, or read the exact value in the box; anything louder than this is treated as clipped and rebuilt. The positive and negative thresholds move together. (–30 to 0 dB, default –0.2 dB.)
- **Learn** — estimates the clip level from the current selection and sets the threshold for you. If the material doesn't show a clipping signature, the threshold is left unchanged.

- **Quality** — how the peaks are rebuilt: **Low** bridges each clipped run with a straight line (fast, no overshoot), **Medium** draws a smooth curve from the surrounding slopes (default), and **High** models the waveform's own oscillation through the gap with an autoregressive reconstruction, best for tonal and bass content (slower). Every tier is clip-consistent — a rebuilt peak is never allowed to dip back below the clip line it was known to have crossed.
- **Makeup gain** — boost applied to the repaired audio to restore level lost to clipping. (0 to 12 dB, default 0 dB.)
- **Post-limiter** — catches any new peaks created by the reconstruction so the output stays below 0 dB and doesn't simply re-clip on export. (On by default.)

How to use it


1. Select the clipped passage (or work on the whole file).
2. Click **Learn**, or drag the Threshold slider just below the flattened peak level.
3. Pick a **Quality** and dial in **Makeup gain** if peaks dropped in level.
4. Use **Preview** to audition the rebuilt peaks.
5. Click **Process** to commit, or **Add to Module Chain** to stack it with other repairs.

Tips

- Set the threshold just under the clip line, not far below it, so you only rebuild genuine clipping and leave clean peaks alone.
- Leave Post-limiter on when adding makeup gain, otherwise the boosted peaks can clip again.
- Start with High quality on exposed transients; drop to Medium or Low if you need speed on long files.

Breath Control

Tame the gasps between words without touching a single syllable.



What it does. Breath Control finds the breaths a speaker takes between phrases and turns them down — or removes them entirely — while leaving the actual speech untouched. It works by spotting short stretches of audio that sit clearly below the talking level, are noise-like rather than tonal, and are about the length of a real breath. Each detected breath is ducked with soft, click-free edges so the result sounds natural, not chopped. It runs entirely on your computer and needs no model download.

When to use it. Reach for it on spoken-word material — podcasts, voiceover, dialogue, audiobooks — whenever loud or distracting inhales pull focus from the words. It is a faster, gentler alternative to hunting down and editing every breath by hand.

Controls

- **Sensitivity** — how eagerly breaths are detected. Higher values catch quieter, softer breaths that sit closer to the speech level; lower values target only the most obvious ones, reducing the risk of clipping real words. (0–1, default 0.5)
- **Reduction** — how much each detected breath is turned down. Values nearer 0 dB leave a subtle, natural breath in place; more negative values duck harder, and –60 dB removes the breath completely. (–60 to 0 dB, default –18 dB)

How to use it

1. Select the speech you want to clean, or leave nothing selected to process the whole file.


2. Set **Sensitivity** so the right breaths are caught, then dial in **Reduction** for how aggressive you want to be.
3. Click **Preview** to hear it in context — if words or consonants start to dip, lower Sensitivity or ease off Reduction.
4. Click **Process** to commit, or **Add to Module Chain** to stack it with other repairs.

Tips

- Start at the default -18 dB. Fully removing breaths often sounds unnaturally clinical; a gentle duck usually reads as cleaner and more professional.
- The detector needs real dynamics between speech and silence, so it has little to do on flat, heavily compressed, or constant-level material.
- Push Sensitivity too high and quiet words or trailing consonants can be mistaken for breaths; back off if speech starts to dip.

De-plosive

Tame the low-frequency thump of p and b pops without thinning the whole voice.



What it does. De-plosive removes the brief low-frequency bursts a mouth makes when it hits the microphone on "p" and "b" sounds. Instead of high-passing everything (which would also drain weight from the whole voice), it listens for the sudden low-end transient that marks a pop and only ducks the low band at that instant. Sustained bass and steady low-frequency program pass through untouched, so the body of the voice stays intact. The reduction fades in fast and out slowly, so the pop disappears cleanly without leaving a click.

When to use it. Reach for it on dialogue, voiceover, and podcast recordings where the talker has popped the mic, or where a missing pop filter has left thumpy plosives scattered through a take. Run it on the whole file, or select just the affected words.

Controls

- **Cutoff** — sets the frequency below which energy is ducked when a plosive is detected. Lower keeps the reduction focused on the deepest thump; higher reaches further up into the low-mids. (80–500 Hz, default 200 Hz.)

- **Threshold** — how loud a low-frequency transient must be to count as a plosive. Lower catches softer pops; raise it if normal speech bass is getting ducked by mistake. (–60 to –10 dB, default –30 dB.)
- **Amount** — how strongly each detected plosive is reduced. At 0 nothing is removed; at 1 the pop is suppressed as fully as the engine allows. (0–1, default 0.85.)

How to use it


1. Select the words containing pops, or leave nothing selected to process the whole file.
2. Set **Cutoff** to cover the thump, then dial **Threshold** so only real pops trigger.
3. Set **Amount** for the depth of reduction you want.
4. Click **Preview** to audition — the pops should soften while the voice stays intact. If words start losing weight, raise **Threshold**.
5. Click **Process** to commit, or **Add to Module Chain** to stack it with other repairs.

Tips

- If the voice starts thinning out, the detector is too aggressive — raise Threshold first, then lower Amount.
- If a heavy pop sounds thin after treatment, nudge Cutoff down rather than pushing Amount higher.
- Because detection keys on transients, you can safely run it over full takes without dulling sustained bass notes or musical low end.

De-ess

Tames harsh sibilance with split-band dynamics.



What it does. De-ess softens the harsh "ess", "sh" and "t" sounds that make close-miked voices sound spitty or sharp. It splits the audio into three bands and applies gentle, automatic gain reduction only to the sibilance band, so the body of the voice and the natural air above it pass through untouched. The reduction is dynamic: it reacts to each sibilant as it happens and lets everything else through cleanly. A live Sibilance Band Response curve sits at the top of the panel, showing exactly which frequencies are ducked and how deeply as you adjust the controls.

When to use it. Reach for it whenever vocals, voiceover or dialogue have piercing or fatiguing sibilance, especially after EQ or compression has pushed the highs forward. It is equally useful for taming bright lavalier mics in post.

Controls

- **Sibilance Band Response** — a log-frequency display (the trough shows where and how much the de-esser dips). Read-only; it tracks the controls below.
- **Threshold** — the level a sibilant must exceed before it is reduced. Lower catches softer, more frequent sibilants; higher acts only on the loudest peaks (-60 to 0 dB, default -28 dB).
- **Cutoff frequency** — the crossover between speech (preserved) and sibilance (reduced). Lower targets lower-pitched "esses"; higher confines the effect to brighter sibilance (2,000 to 12,000 Hz, default 4,500 Hz).

- **Amount** — how strongly detected sibilance is reduced. Higher means deeper ducking (0 to 100%, default 80%).
- **Speed** — *Fast* for crisp, transient sibilants; *Slow* for smoother, more transparent gain riding (default *Fast*).

How to use it

1. Select the vocal region (or leave nothing selected to process the whole file).
2. Set the **Cutoff** so the trough sits over the offending sibilance, then dial **Threshold** and **Amount** to taste.
3. Use **Preview** to A/B in context, then **Process** to commit. Add it to your Module Chain to reuse the settings.

Tips

- If the vocal starts losing vowels or breath body, raise the Cutoff or Threshold so only the sibilance is caught.
- Choose *Slow* on sustained or operatic vocals to avoid a lispy artifact; keep *Fast* for spoken word.
- Over-de-essing dulls a voice — back off Amount until the harshness is gone but the consonants stay intelligible.

Presets


- **Gentle Vocal** — light, transparent taming for already-clean vocals.
- **Vocal De-ess** — the everyday starting point for sung or spoken vocals.
- **Harsh Sibilance** — aggressive reduction for spitty, over-bright recordings.
- **Broadcast** — controlled, consistent sibilance for voiceover and dialogue.

Utility Modules

Everyday signal utilities — equalisation, generation, channel and stereo handling, phase and sample-rate conversion.

EQ Filters

Shape your tone with a hands-on parametric equalizer.



What it does. EQ Filters is a parametric equalizer that lets you boost, cut, or filter specific frequency ranges. It opens with four bands and a large interactive response display: a filled curve shows the combined shape your bands are making, and each enabled band appears as a coloured dot you can drag. You can also dial in exact values in the band editor and the band list below. While audio plays, a live spectrum analyzer draws the output's spectrum behind the curve, so you can watch your changes land on the actual material. The curve is purely a filter — it adds nothing of its own, so disabled or flat bands have no effect on the sound.

When to use it. Reach for it to tame harshness, add air, roll off rumble, thin out a muddy low end, or carve a notch around a resonant ring. It works on dialogue, music, and full mixes alike, and pairs well with corrective work earlier in your chain.

Controls

- **Response display** — the large plot at the top. Drag a band's dot left/right to change its **frequency** and up/down to change its **gain**. High-pass and low-pass bands only move sideways (they have no gain). Click a band to select it; the selected dot grows slightly.
- **Band editor** (appears when a band is selected) — a card with the precise controls for that band:
 - **Band [shape]** — the filter shape: Low Shelf, Peak, High Shelf, High Pass, or Low Pass.
 - **Freq** — the centre frequency the band acts on, or the cutoff for high/low-pass (20 Hz–20 kHz).
 - **Gain** — boost (+) or cut (–) at that frequency, from –24 to +24 dB (default 0). Hidden for high/low-pass bands.
 - **Q** — bandwidth; higher Q makes a narrower, more focused band (0.1–10). Hidden for high/low-pass bands.
 - **Slope** — the roll-off steepness of a High Pass or Low Pass band: **12 dB/oct**, **24 dB/oct**, **48 dB/oct**, or **Brickwall** (a near-vertical cut). Shown only for pass bands (in place of Gain and Q); the steeper options are maximally-flat Butterworth cascades, so a sharper slope adds no resonant bump at the cutoff. Defaults to 12 dB/oct.
- **Band list** — one row per band showing its colour, number, shape, frequency, and gain. Click the dot at the left of a row to enable or bypass that band without removing it; click the row to select and edit it.
- **Add Band** — appends a fresh peak band at 1 kHz (or click an empty spot on the curve to drop one there); remove a band again from the trash button in its editor — when you delete a band the editor stays open on the band that slides into its place (the next one down, or the new last row). At least one band always remains, so the trash button is disabled on the final band.
- **Reset** — restores all bands to their flat default settings.
- **Phase** — chooses the EQ's filtering style: **Linear (FIR)** or **Natural (IIR)**. Linear phase (the default) shapes only the tone, so a boost or cut never rings — there is no phase-shift resonance; it adds latency, which this offline editor compensates for automatically. Natural phase is zero-latency minimum-phase filtering. This is a global setting: it applies to every EQ instance, including Module Chain steps, so it appears only in the live editor and not in the chain/batch settings editor.
- **Analyzer** — shows the live output spectrum behind the EQ curve while audio plays, Pro-Q style (on by default). Turn it off for a cleaner plot; the choice is remembered across sessions. The toggle lives only in this panel — the chain/batch settings editor has no live playback, so it is omitted there.
- **Presets** — the standard module preset row above the panel carries the EQ's built-in profiles as read-only **Factory** presets: *Flat / Reset*, *Low Cut*, *De-Mud (low-mid cut)*, *Vocal Presence*, *Bass Boost*, *Air / Brightness*, and *Telephone*. Pick one to drop its band set straight in, then tweak and save your own alongside them as user presets. (The old in-panel preset menu was retired; any presets you had saved there are folded into the row automatically the first time you open it.)

How to use it

1. Select the region to process, or leave nothing selected to treat the whole file.
2. Drag dots in the plot for a quick shape, then click a band and fine-tune **Freq**, **Gain**, and **Q** in the editor.
3. Toggle bands on and off from the list to compare.


4. Use **Preview** to audition — it is a non-destructive listen toggle (Shift+Space): press it to render and hear the change, press again to stop the audio without losing the rendered preview. When it sounds right, hit **Process** to commit. Choose **Add to Module Chain** to stack the EQ with other modules.

Tips

- Leave **Analyzer** on and keep the audio playing while you drag bands — the spectrum updates behind the curve, so you hear *and* see each move land in real time.
- **High Pass** with a gentle **12 dB/oct** slope is the cleanest way to remove low rumble without thinning the body; reach for a steeper 24 or 48 dB/oct — or **Brickwall** — only when a hard cut is really needed.
- Leave **Phase** on **Linear (FIR)** for mastering-grade tonal moves where you want a boost or cut with no ringing; switch to **Natural (IIR)** if you prefer classic minimum-phase behaviour or want zero added latency.
- For surgical fixes, push Q high and a Peak band's gain down, then sweep the frequency to find the offending ring before cutting.
- Gentle, wide shelves (low Q) sound the most natural for broad tonal balancing.
- Bypassed bands are remembered, so you can keep a setting parked and bring it back instantly.

Signal Generator

Synthesize tones, sweeps, impulses, silence and shaped noise, then drop them in at the cursor or over a selection.



What it does. The Signal Generator builds audio from scratch — pure tones, frequency sweeps, single-sample impulses, silence and several flavours of noise — and writes it into your document. The generated clip matches your file's sample rate and channel count automatically. You decide whether it replaces a selection, mixes on top of existing audio, or pushes everything aside and inserts.

When to use it. Lay down a reference tone for calibration, drop a log sweep for room or gear measurement, build a test-noise bed, insert a gap of clean silence, or create impulse markers for timing checks.

Controls

- **Signal type** (tab bar) — choose **Tone**, **Sweep**, **Noise**, **Impulse** or **Silence**. The controls below change to match.
- **Waveform** (Tone, Sweep) — **Sine**, **Square**, **Triangle** or **Sawtooth** (default Sine).

- **Frequency** (Tone) — pitch of the tone (1–22000 Hz, default 440).
- **Pulse Width** (Tone, Square only) — duty cycle of the square wave; lower is thinner, higher is fatter (0.05–0.95, default 0.50).
- **Direction** (Sweep) — sweep **Up** or **Down**.
- **Scale** (Sweep) — **Log** (musically even, default) or **Linear**.
- **Start / End** (Sweep) — sweep's beginning and ending frequency (1–22000 Hz; defaults 440 / 1000).
- **Noise Type** (Noise) — **White Gaussian** (default), **White Uniform**, **White Triangular**, **White Binary**, **Pink**, **Brown** or **Blue**.
- **Impulse Rate** (Impulse) — clicks per second (0.1–200 Hz, default 440).
- **Paste Mode** — **Replace** (overwrite the selection, or insert at the cursor), **Mix** (sum onto existing audio), or **Insert** (push later audio rightward).
- **Channels** — **All Channels**, **Left Only**, **Right Only**, **Alternating Polarity** or **Mirrored Sweep**.
- **Duration** — clip length (0.001–300 s, default 1). **Use Selection** sets this to the current selection's length.
- **Peak / RMS Level** — output level; labelled RMS for noise, Peak otherwise (–120 to 0 dBFS, default –12). Hidden for Silence.
- **Fade In / Fade Out** — edge ramps to avoid clicks (0–5000 ms, default 2 each).

How to use it

1. Pick a **Signal type** and dial in its parameters.
2. Set **Duration**, or click **Use Selection** to match a highlighted range.
3. Set **Level**, **Fade In/Out** and **Channels**.
4. Choose a **Paste Mode**, then place your cursor or make a selection.
5. Click **Render** to write the signal into the document.

Tips

- With nothing selected, Render lands at the cursor; with a selection, Replace overwrites it and Mix layers onto it.
- Keep small fades on tones and noise — they prevent edge clicks at the seams.
- Use a Log sweep for measurement work; it spends equal time per octave.
- Channels options like Mirrored Sweep and Alternating Polarity are handy for stereo and phase tests.

Phase

Flip polarity, rotate phase, and time-align channels — surgical fixes for cancellation and stereo smear.



What it does. Phase gives you three related tools in one panel. It can flip a channel's polarity (turn the waveform upside down), rotate the phase of the whole signal by any angle without changing its loudness or magnitude, and nudge one channel forward or backward by a handful of samples to time-align a stereo pair. A live before/after scope shows a dim "ref" trace and a bright trace so you can see exactly what your settings do to the wave.

When to use it. Reach for Phase when a stereo recording sounds hollow or collapses in mono (a polarity or alignment problem), when a double-miked source is partly cancelling, or when an asymmetric waveform (speech, brass, kick) clips on one side and you want to tame its peak without losing level.

Controls

- **Phase Shift Scope** — a two-trace oscilloscope over the synthetic reference cycle. The ghost trace is the original; the glowing trace follows your Rotation and Delay. It turns to a warning colour and reads "INVERTED" at exactly $\pm 180^\circ$, and shows a direction caret when a delay is set.
- **Polarity target** — tabs choosing which channels Invert affects: **Selected**, **All**, **L**, **R** (mono files show Sel / All / L). Default **Sel**.
- **Invert Polarity** — flips the chosen channels upside down (equivalent to a 180° shift). Disabled when no valid channel is targeted.
- **Rotation** — rotates phase across the whole spectrum, -180° to $+180^\circ$ (default **0°**); $\pm 180^\circ$ equals a polarity flip. **Suggest** finds the angle that lowers the peak the most; **Reset** returns it to 0° .
- **Delay target** — tabs choosing the channel to time-shift: **L** / **R** (mono: L only). Default **R**.
- **Samples** — inter-channel shift, -512 to $+512$ samples (default **0**); type an exact value or drag. **Apply Delay** commits it.
- **Alignment — Detected** — read-out of the measured L/R offset. **Detect** measures the offset without changing audio; **Auto Align** time-shifts the right channel to match the left (and inverts it if needed). Both need a stereo file.

How to use it


1. Select the region (or whole file) you want to treat.
2. For a quick stereo fix, click **Detect** to see the offset, then **Auto Align**.
3. For peaks, click **Suggest**, then **Preview** to compare, and **Apply** to commit the rotation.
4. To fix cancellation, set the **Polarity target** and click **Invert Polarity**; to dial in a manual offset, set **Samples** and **Apply Delay**.
5. Add the module to a **Module Chain** to reuse the same rotation and delay across files.

Tips

- Rotation changes peak shape but not loudness — it's a clip-safe alternative to pulling gain down.
- If Detect reports "inv", the channels are out of polarity; Auto Align will flip and align in one move.
- Polarity and Delay buttons act immediately on the audio; Rotation flows through Preview/Apply, so audition before committing.
- Keep Delay small (a few samples) — large shifts are for true alignment, not tone.

Azimuth

Re-square a stereo image skewed by misaligned tape heads.



What it does. Azimuth corrects inter-channel time and gain skew, the classic problem of a tape machine whose head was tilted slightly out of alignment so the two channels no longer arrive level or in step. It lets you nudge each channel's level independently and slide the right channel forward or backward in time until the two sides line up again. A Stereo Alignment Field at the top of the panel pictures both channels as stacked rays: each ray's thickness and glow track its gain, the right ray slides sideways with the delay, a dashed tie-line shows the gap being closed, and a balance pointer leans toward the louder side. The picture is drawn from your settings, not from a live analysis of the audio.

When to use it. Reach for Azimuth on digitised reel-to-reel or cassette transfers where the highs sound smeared, the image drifts off-centre, or one channel is consistently softer than the other. It is also handy any time a stereo pair is slightly out of level or time-aligned.

Controls

- **Stereo Alignment Field** — the visualization described above. It updates instantly as you move the sliders and shows a live readout of L/R gain in dB and the right-channel delay in both samples and milliseconds.
- **L Gain** — level applied to the left channel; lower cuts, higher boosts. (–24 to +24 dB, default 0.)

- **R Gain** — level applied to the right channel; lower cuts, higher boosts. (–24 to +24 dB, default 0.)
- **R Delay** — shifts the right channel in time to realign it with the left; negative is earlier, positive is later. (–512 to +512 samples, default 0.)

How to use it


1. Select the passage to correct (or work on the whole file).
2. Set **L Gain**, **R Gain** and **R Delay** while watching the alignment field.
3. Press **Process** in the footer to commit.
4. To stack Azimuth with other repairs, use the footer **+** menu's **Add to Module Chain**.

Tips

- Small delays go a long way: tape azimuth error is usually a handful of samples, so reach for fine R Delay nudges before large ones.
- Process and the Module Chain do nothing at identity (zero gain and zero delay), so set a real correction before committing.

Center Extract

Pull the vocal out of the middle of a stereo mix — or push it away.



What it does. Center Extract works on the phantom center of a stereo file — the lead vocal, kick, snare and bass that sit dead-centre because they're equal and in-phase in both channels. In **Extract** and **Remove** it works in the frequency domain: for every time-frequency bin it measures how equal-and-in-phase the left and right spectra are (a coherence mask), then either keeps only the shared centre content (Extract) or keeps only what differs between the channels (Remove, the classic karaoke trick). Because it's per-bin rather than a blunt channel sum, a hard-panned instrument or one-sided noise stays where it belongs instead of folding into the result. The third mode, **Balance**, is a plainer broadband mid/side re-gain that just re-levels the centre and the sides. It needs a stereo file — the apply buttons are disabled on mono.

When to use it. Reach for it to make a quick karaoke or backing track, to pull a phantom-centre vocal forward (or duck it under the music), or to widen a mix by lifting the sides relative to the centre.

Controls

- **Stereo field display** — a top-down schematic of the image: a bright central beam for the phantom centre and two fanning side wedges. In Extract the centre beam narrows as you raise Strength; in

Remove the centre beam collapses to a hollow dashed mark. In Balance the beam and wedge heights track the centre and side gains, with a live readout of the resulting centre/side levels in dB and an overall width percentage.

- **Mode tabs** — **Extract** keeps only the phantom-centre content; **Remove** keeps everything but the centre (karaoke); **Balance** re-levels the centre and the sides by hand. Extract and Remove share a single **Strength** control; Balance swaps it for the **Center** and **Sides** gain sliders.
- **Strength** (*Extract / Remove*) — how aggressively off-centre (panned or uncorrelated) content is rejected by the coherence mask. Higher isolates the shared phantom centre more surgically; lower is gentler (0–100 %, default 75 %).
- **Center** (*Balance only*) — gain applied to the phantom-centre content. Lower it to push vocals back, raise to bring them forward (–48 to +12 dB, default 0).
- **Sides** (*Balance only*) — gain applied to the stereo sides. Lower to narrow the image, raise to widen it (–48 to +12 dB, default –12).
- **Analyze** — measures the current selection and reports Center level, Sides level and the centre-to-sides ratio (C/S), so you can judge what's there before committing.

How to use it

1. Select the stereo region you want to treat (no selection processes the whole file).
2. Click **Analyze** to read the centre/sides balance.
3. Pick a mode: Extract, Remove or Balance.
4. In Extract or Remove, set **Strength** for how surgically off-centre content is rejected. In Balance, set the **Center** and **Sides** levels while watching the field display.
5. Click the apply button (**Extract Center**, **Remove Center** or **Rebalance**) to process every stereo pair. The analysis refreshes so you can confirm the result.

Tips


- Karaoke vocals never vanish completely — anything panned or with stereo reverb survives, so judge by ear.
- Use the readout's width percentage as a quick mono/stereo gauge while you dial in Balance.
- The analysis re-measures after edits and undo; if it looks stale, run **Analyze** again.

Presets

- **Vocal Boost** — Balance mode: lifts the centre and trims the sides to bring a lead vocal forward.
- **Vocal Reduce (Rebalance)** — Balance mode: deeply ducks the centre for a backing track without full removal.
- **Widen Sides** — Balance mode: raises the sides for a broader image while leaving the centre untouched.
- **Isolate Center** — Extract mode at high Strength, to leave only the phantom-centre content.
- **Remove Center (Karaoke)** — Remove mode at high Strength, to strip the centred vocal for a backing track.

Resample

Mastering-grade sample-rate conversion — or a simple rate relabel.



What it does. Resample changes your file's sample rate using a high-quality windowed-sinc converter with a flat passband and deep alias rejection, keeping length exactly correct. A status chip at the top shows the conversion at a glance (for example "44100 → 48000 Hz · upsample") so you always know which way you are going, and a live response graph shows exactly what the anti-alias filter is doing. You can also choose to *relabel* the rate without touching the audio, and — when upsampling — optionally rebuild the missing high frequencies.

When to use it. Use it to deliver a file at a specific rate (44.1 kHz for music, 48 kHz for video/post), to consolidate mismatched material into one project rate, or to fix a file whose declared rate is simply wrong.

Controls

- **New rate** — the target sample rate in Hz. Type any value (clamped between 1,000 and 768,000 Hz) or pick a common rate from the chevron menu (8,000 up to 192,000 Hz). Default 48,000 Hz.

- **Filter response graph** — a plot of the anti-alias response around the target Nyquist: a dashed **Ideal** brick wall, the **FFT** curve (the windowed-sinc filter actually applied), and the shaded **Aliasing** region — the filter's residue folded back below Nyquist. It redraws live as you move **Steepness** and **Cutoff**, so you can watch the roll-off sharpen and the aliasing sink. Hidden when relabelling.
- **Steepness** — filter sharpness. Higher gives a steeper cutoff and better alias rejection but more time-domain ringing; lower is gentler (0.00–1.00, default 0.60). Hidden when "Change tag only" is on.
- **Cutoff** — low-pass point as a fraction of the Nyquist limit. Lower widens alias rejection at the cost of top-end passband; higher keeps more highs (0.800–1.000, default 0.920). Hidden when relabelling.
- **Post-limiter (≤ -0.3 dB)** — catches inter-sample/true-peak overshoots the converter can introduce, keeping the result under the ceiling (default off).
- **Change tag only (relabel)** — changes the declared rate *without* resampling. This is a metadata fix that alters playback speed and pitch; use it only to correct a mislabelled file (default off).
- **Rebuild highs** — appears only when upsampling. After conversion it regenerates the band above the original Nyquist with an on-device AI super-resolution engine (default off).
- **Algorithm** (shown when Rebuild highs is on) — picks what fills the new highs, from one menu ordered fastest to slowest: **Fast (LavaSR)**, **Harmonic** (model-free, always available), **Neural Speech (AERO)**, **Neural Music (AERO)**, and **Studio (MossFormer)** = MossFormer2-SR (the default). It is the same menu as the Spectral Recovery module. If the chosen algorithm's model isn't installed, a greyed notice tells you what will happen instead: the AERO neural algorithms fall back to the model-free Harmonic rebuild, while Fast/Studio leave the upsampling plain DSP until you install the model.

How to use it

1. Set **New rate** by typing or picking from the menu.
2. For a true conversion, leave "Change tag only" off and set **Steepness** and **Cutoff**, watching the response graph as you go; enable **Post-limiter** if peaks concern you.
3. When upsampling, optionally turn on **Rebuild highs** and pick the **Algorithm** that matches your material.
4. Click **Process** to apply. Note that Resample cannot be added to the Module Chain — it changes the sample rate, so run it on its own.

Resampling rewrites the whole file as one undoable edit, and markers and regions are rescaled to track the new rate.


Tips

- For pristine downsampling, lower the **Cutoff** slightly to push aliasing further below the noise floor — you can see the shaded aliasing curve drop in the graph as you do.
- Reach for **Change tag only** only when the speed/pitch shift is exactly what you want — it does not convert.
- **Rebuild highs** is best on music and sound effects; it cannot truly recreate detail that was never recorded.

- Match the rebuild **Algorithm** to the material: **Neural Speech (AERO)** for voice, **Neural Music (AERO)** for instruments and mixes, **Harmonic** or **Fast (LavaSR)** when speed matters, and **Studio (MossFormer)** as the heavyweight general-purpose default.

Stereo Width

Widen, narrow, or build a stereo image — while keeping the low end solid.



What it does. Stereo Width reshapes the left/right image of your audio. On a stereo file it spreads or narrows the sides — pulling everything to dead-centre mono at one extreme, or exaggerating the spread at the other — without altering what your mix sounds like when summed to mono. It can also lock the bass below a chosen crossover to mono so the low end stays phase-solid. On a mono file the panel switches automatically and offers to synthesize a believable stereo image from the single channel.

When to use it. Open up a narrow stereo recording, rein in an over-wide mix, keep kick and bass centred for vinyl or club playback, or give a mono voice or instrument a sense of space.

Controls

- **Stereo Field** — a goniometer-style display that previews the setting: a thin vertical bar at 0% (mono), a diamond at 100%, and a wider wedge above that. It reflects the controls, not live audio.
- **Width** — how wide the image sits. 0% collapses to mono, 100% leaves it unchanged, 200% exaggerates the spread (0–200%, default 100%).
- **Bass mono below crossover** — when on, sums everything under the crossover to mono so the low end stays phase-coherent (default off).
- **Crossover** — the frequency below which the bass is folded to mono. Active only when bass-mono is on (40–500 Hz, default 120 Hz).

- **Stereo Field readout** — live **Correlation** and **Width** figures for the current selection, so you can judge the image numerically.
- **Apply Width** — bakes the width and bass-mono settings into the audio.

For mono files instead: an **Amount** slider (0–1, default 1) controls how much synthesized width to create, with a **Convert to Stereo** button.

How to use it

1. Select the region to treat (or leave unselected for the whole file).
2. Set **Width**, and enable **Bass mono** with a **Crossover** if you want a solid low end.
3. Watch the Stereo Field display and the Correlation/Width readout.
4. Click **Apply Width** (or **Convert to Stereo** on a mono file) to commit. Note that Stereo Width cannot be added to the Module Chain — it can change the channel count, so run it on its own.

Tips


- Keep an eye on **Correlation**: values near +1 are very mono, near 0 are wide, and negative values warn of phase issues that can hollow out on mono playback.
- For masters destined for vinyl or club systems, enable bass-mono around 100–150 Hz to keep the kick and bass tight.
- Pushing **Width** above 100% only amplifies what is already there — it cannot add space to a truly mono source; use **Convert to Stereo** for that.

Presets

- **Mono** — collapses the image fully to centre.
- **Natural** — leaves the width unchanged.
- **Wide** — a moderate widening past unity.
- **Wide + Mono Bass** — extra width up top with everything below 120 Hz kept mono.

Channels

Reshape your file's channel layout — fold to mono, open up to stereo, pan, fix polarity, and work in mid/side.



What it does. Channels is your toolbox for everything to do with channel layout and routing. From one panel you can fold a multichannel or stereo file down to mono, build a stereo pair from a mono recording, swap a reversed left/right pair, pan or balance the image with constant-power law, flip polarity to repair an out-of-phase channel, and switch into mid/side to widen the image or edit the Mid and Side signals directly. It is multichannel aware: a 5.1 or 7.1 file gets a proper surround-to-stereo fold, while other channel counts are folded sensibly. Operations rewrite the whole file, so they change its channel count rather than editing a region.

When to use it. Reach for it whenever the channel format is wrong for the job: collapsing an interview to mono for a podcast feed, giving a mono voiceover a stereo footprint, rescuing a track with one channel wired backwards, or narrowing a too-wide mix before mastering.

Controls

- **Operation** — picks the task: **To Mono**, **To Stereo**, **Swap L / R**, **Pan**, **Phase Fix**, or **Mid/Side** (default To Mono). The panel below changes to match. A status chip at the top shows your file's current layout (mono, stereo, or multichannel), and a small note appears when the chosen operation would do nothing.
- **To Mono source** — **Mix** averages every channel together, **Left** keeps only the left channel, **Right** keeps only the right (default Mix).
- **To Stereo source** (mono files only) — **Duplicate** copies the mono signal to both sides, **Pan** places it using the Pan slider (default Duplicate). Stereo and wider files fold automatically instead.

- **Swap L / R** — exchanges the left and right channels. It has no options: choose it and Process. On a mono file the hint reads "Mono file — nothing to swap."; on files wider than stereo it notes "Swaps the first two channels (L / R)." and leaves the remaining channels untouched.
- **Pan** — constant-power pan/balance from -100 (hard left) through 0 (centre) to +100 (hard right) (default 0). On a mono file it positions the signal; on stereo it balances without ever boosting above unity.
- **Phase Fix target** — inverts polarity for **All** channels, or just **Left** or **Right** on a stereo/wider file (default All).
- **Mid/Side mode** — **Width** widens or narrows the image, **Encode** exposes the raw Mid and Side signals for direct editing, **Decode** restores them back to left/right (default Width). Needs a stereo file.
- **Width** (Mid/Side, Width mode) — 0% collapses to mono, 100% leaves the image unchanged, 200% doubles the width (range 0–200%, default 100%).

How to use it

1. Open the Channels module and check the layout chip so you know what you're starting from.
2. Choose an **Operation**, then set its options (source, pan, target, or width).
3. Watch for the grey hint line — if it says the operation is a no-op for this file, adjust your choice.
4. Click **Process** to apply it to the whole file. Note that Channels cannot be added to the Module Chain — it can change the channel count, so run it on its own.

Tips


- Before mixing to mono, audition each channel: if **Mix** sounds hollow, one channel may be out of phase — run **Phase Fix** first, then fold.
- Use **Mid/Side → Encode** to edit the centre and the sides separately (for example, clean up reverb in the Side), then **Decode** to put the file back to normal stereo.
- The **Width** control at 0% is a quick, phase-safe way to mono-check a mix without leaving the module.
- Pan/balance never adds gain, so it won't introduce clipping — but folding a hot surround mix to stereo is automatically kept clip-safe for you.

Process Modules

Creative and corrective processors — dynamics, reverb, fades, gain and loudness, time and pitch, and reference matching.

Distance

Move a source toward or away from the mic — pull closer or push it back into the room.



What it does. Distance is a single bipolar control that changes how near or far a sound seems. Pull toward **Closer** and FFT dereverberates the source — lifting it out of the room — then restores the intimacy of a close mic (fuller lows, more presence, sharper transient bite), all loudness-matched so you hear a change in distance, not a jump in level. Push toward **Farther** and it does the opposite: it sets the source into a synthetic room, lowers the direct level, gently rolls off the top, and thins the close-mic bass. It works on any content — voice, instruments, sound effects — using signal processing only, with no model to download.

When to use it. Tame an over-reverberant dialogue or instrument recording and bring it forward in the mix; or take a dry, in-your-face source and place it convincingly back in a space. Useful for matching the perceived distance of two clips that were recorded differently.

Controls

- **Distance** — the one bipolar control. Drag left toward **Closer** to dereverberate and restore proximity; drag right toward **Farther** to add room and push the source back. **0** in the centre is bypass (no change). The further you go in either direction, the stronger the effect. A status line above the control reads *Pulling closer, Pushing away, or Bypass (centre)* so you always know which way you're working. (Range -1 to +1, default 0.)

How to use it

1. Select the region you want to move (or work on the whole file).
2. Drag **Distance** toward Closer or Farther.

3. Click **Preview** to audition the result; nudge the control until the placement feels right.
4. Click **Process** to apply it, or **Add to Module Chain** to stack it with other modules.

Tips

- Pull-closer is gentle by design on sustained, dense material — it dereverbs gappy, transient sources (speech, plucks, percussion) most cleanly and avoids over-drying instruments that ring naturally.
- Because the closer side is loudness-matched, judge it by space and tone, not volume.
- A small amount usually convinces; extreme settings on the Farther side add the most obvious synthetic room.

Reverb

Drop a sound into any space — from a tight booth to a cathedral — with a tail that stays smooth and natural.



What it does. Reverb wraps your audio in a lush, modulated tail that simulates the reflections of a real space. It is fully synthesised, so it works on any material — voice, instruments, drums, full mixes — and the tail is engineered to stay dense and smooth without ever turning metallic or ringy. You shape everything by ear: the size of the room, how long it rings out, its tone, its texture, and how wide it spreads across the stereo field.

When to use it. Add depth and glue to a dry vocal or voiceover, place a close-mic'd instrument in a believable room, or create a long ambient wash for sound design. Reach for it whenever a recording sounds flat, boxed-in, or unnaturally close.

Controls

- **Decay-vs-frequency curve & tail preview** — At the top, drag the three dots to shape the space: **Low** sets the bass bloom, **Mid** sets the overall reverb time, and **High** sets the high-frequency damping. The slim strip below previews the resulting tail. These dots mirror the Low Decay, Decay and Damping sliders.
- **Size** — Scales the whole space from a tight booth to a cathedral. Higher is bigger and more open (0–1, default 0.55).
- **Decay** — How long the mid-range tail rings out. Higher means a longer reverb (0.2–12 s, default 1.9 s).
- **Pre-delay** — A gap before the tail begins, keeping the dry transient clear and separated from the reverb. Higher pushes the tail later (0–250 ms, default 22 ms).
- **Damping** — High-frequency absorption. 0 is bright with long, airy highs; 1 is dark with fast-fading highs (0–1, default 0.5).
- **Low Decay** — Bass decay relative to the mids — warmth and low-end bloom. Higher means a fatter, longer low end (0.25–2.5×, default 1.25×).
- **Diffusion** — Texture of the reflections. Low gives looser, more discrete echoes; high gives a dense, smooth wash (0–1, default 0.7).
- **Modulation** — Gently moves the tail to keep it lush and prevent any metallic ring. Higher is more shimmer and motion (0–1, default 0.4).
- **Mix** — Dry/wet blend. 0 is fully dry, 1 is fully wet (0–1, default 0.30).
- **Width** — Stereo spread of the reverb, mono-compatible at every setting. 0 is centred, 1 is fully wide (0–1, default 1.0).

How to use it

1. Select the region you want to treat, or leave nothing selected to process the whole file.
2. Pick a starting point from the Factory presets menu, or dial in Size and Decay by hand.
3. Drag the Low/Mid/High dots and adjust Damping, Low Decay and Diffusion to taste, watching the tail preview.
4. Set the **Mix** for the dry/wet balance you want, then **Preview** to audition in context.
5. **Process** to commit, or use **Add to Module Chain** to stack it with other modules before printing.

Tips


- Keep **Mix** low (around 0.15–0.30) on lead vocals and dialogue — a little tail adds depth without washing out clarity.
- A longer **Pre-delay** keeps consonants and transients crisp while still giving a big tail; great for vocals.
- If a long, bright tail feels harsh, raise **Damping** or pull the High dot down rather than shortening the whole decay.
- Push **Modulation** higher for lush, evolving ambient pads; lower it for tighter, more static rooms.

Presets

- **Small Room** — Short, tight ambience for a touch of natural space.
- **Ambience** — Subtle room tone that adds air without an obvious tail.
- **Plate** — Bright, dense studio-plate sheen for vocals and snares.
- **Medium Hall** — Balanced all-purpose hall.
- **Large Hall** — Big, open hall with a long, rich tail.
- **Cathedral** — Vast, slow space for dramatic, cinematic decays.
- **Lush Vocal** — Long, modulated, warm tail tuned to flatter lead vocals.

Dynamics

Even out the loud and quiet, then catch the peaks.



What it does. Dynamics is a compressor that tames anything louder than the threshold, with optional makeup gain to bring the level back up, followed by a true-peak safety limiter that never lets the output exceed your ceiling. A live transfer curve plots the current settings as an input-to-output level law, so you can see exactly how loud a given input becomes before you commit. The compressor and limiter both react to the loudest channel, so stereo placement stays intact.

When to use it. Reach for it to level an uneven vocal or voiceover, glue a drum bus or mix together, or rein in stray peaks. The always-on limiter also makes it a safe final stage to guarantee nothing clips.

Controls

- **Transfer curve** — the graph at the top. The faint diagonal is unity (no change); the bright line is your current compression and limiting law. A vertical marker shows the threshold, and a readout reports the in/out level at the threshold. Watch it flatten as you raise the ratio or lower the ceiling.

- **Threshold** — the level above which compression starts; lower values compress more of the signal (–60 to 0 dB, default –18 dB).
- **Ratio** — how strongly levels above the threshold are reduced; higher is stronger (1:1 to 20:1, default 4:1).
- **Attack** — how fast compression engages. Short attacks catch transients, longer ones let them punch through (0.1 to 200 ms, default 5 ms).
- **Release** — how fast gain recovers after the level falls below the threshold (5 to 1000 ms, default 100 ms).
- **Makeup** — output gain added after compression to restore loudness (0 to 24 dB, default 0 dB).
- **Limiter Ceiling** — the hard output ceiling applied after the compressor, always active (–12 to 0 dB, default –0.3 dB).

How to use it

1. Select a region, or leave the selection empty to process the whole file.
2. Set the **Threshold** and **Ratio**, watching the transfer curve.
3. Shape the timing with **Attack** and **Release**, then add **Makeup** to taste.
4. Confirm the **Limiter Ceiling** is where you want it.
5. **Preview** to listen, then **Process** to apply. You can also **Add to Module Chain** to stack it with other processing.

Tips


- Set a high ratio with a low threshold to use Dynamics as a standalone brickwall limiter.
- The ceiling is always on, so even at extreme settings your output stays clean of peaks.
- Add Makeup gradually and compare against the unprocessed signal so you match loudness, not just push it louder.
- Long releases give smooth "glue"; short releases sound more aggressive and can pump.

Presets

- **Vocal Leveler** — smooths an uneven vocal or voiceover.
- **Drum Bus** — punchy, faster control for drums.
- **Gentle Glue** — soft, slow compression to bind a mix together.
- **Brickwall Limiter** — high ratio, fast attack for hard peak limiting.

Sound Enhance

Four friendly sliders that simply make the audio sound better.



What it does. Sound Enhance is a one-stop sweetening pass: it adds finish and smoothness, fills out the low end, opens up the top, and sharpens attacks — each from its own slider, with no engineering vocabulary in between. Under the hood it keeps the output level under control with a built-in safety ceiling, so piling on enhancement never pushes the result into clipping.

When to use it. When a recording sounds flat, dull, thin or lifeless and you want a quick, musical lift without building an EQ + compressor + exciter chain by hand. It suits podcasts, voice-overs, demos and rough mixes that need to sound presentable fast.

Controls

- **Display** — a spectral curve across the top shows the combined enhancement shape, with four colored bars labeled **Warm**, **Polish**, **Punch** and **Shine** that rise with their sliders.
- **Polish** — overall finish and smoothness (0–1, default 0.45).
- **Warmth** — fuller low and low-mid tone (0–1, default 0.30).
- **Shine** — open top-end gloss (0–1, default 0.35).
- **Punch** — more snap on attacks (0–1, default 0.30).

How to use it

1. Select the audio to sweeten (or leave nothing selected to treat the whole file).


2. Click **Preview** and move the four sliders while listening — there is no Learn step, so you can start adjusting immediately.
3. Click **Process** to commit, or **Add to Module Chain** to stack it with other processing.

Tips

- Start from the defaults and adjust one slider at a time — the controls interact, and small moves usually sound the most natural.
- With all four sliders at 0 the module is a bypass, so Preview and Process stay disabled until you raise something.
- The internal ceiling keeps peaks safe, but heavy settings still change the character noticeably — A/B with Bypass in the footer.
- Sound Enhance runs in the Module Chain and in Batch, so a setting you like can finish a whole folder of episodes.

Fade

Shape a smooth level ramp into or out of your selection.



What it does. Fade applies a shaped fade-in or fade-out across your selection, ramping the level smoothly between silence and full volume. You choose the direction and the shape of the ramp, and a live envelope graph at the top of the panel draws the exact gain-versus-time curve the fade will apply — the line you see is the effect. A small ring marks the silent end of the curve, and gridlines show where the level passes 0 dB and -6 dB.

When to use it. Use it to ease a clip in from silence or out to nothing without clicks, to clean up an abrupt edit point, or to set up a smooth transition into a crossfade. It's equally at home on a single word of dialogue, a music tail, or the head of a sound effect.

Controls

- **In / Out** — picks the direction. *In* ramps the level up from silence; *Out* ramps it down to silence. The envelope graph updates to match (default: In).
- **Curve** — the shape of the level trajectory: *Linear* (straight ramp), *Logarithmic* (slow start, fast finish), *Equal Power* (constant-energy curve, good for crossfade halves), *S-Curve* (gentle at both ends), and *Cosine* (raised-cosine, smooth for tonal material).

- **Selection / Seconds** — chooses how far the fade spans. *Selection* fades across the whole selected region; *Seconds* fades over a fixed length you set, anchored at the start (fade-in) or end (fade-out) of the selection (default: Selection).
- **Duration** — appears only in *Seconds* mode; sets the fade length in seconds (0.01–10 s, default 0.25 s).

How to use it


1. Select the region you want to fade (or the head/tail of a clip).
2. Choose **In** or **Out** and pick a **Curve**, watching the envelope graph preview the shape.
3. Choose **Selection** to fade the whole region, or **Seconds** and dial in a **Duration**.
4. Press **Process** in the footer to write it now — Process applies the fade in whichever direction you've set. To queue it as a step in your processing chain instead, use the footer **+** menu's **Add to Module Chain**.

Tips

- A fade does nothing on a region that's already silent — give it audio to ramp.
- For matched crossfades between two clips, use *Equal Power* on both halves to keep loudness constant through the join.
- In *Seconds* mode, the fade is pinned to the edge of the selection, so you can fade just the first or last fraction of a long selection.
- Applied fades are undoable like any edit, so audition different curves freely.

Gain

Set the level — by a fixed amount, or to a peak or loudness target.



What it does. Gain changes the overall level of your audio in one of three ways. You can add or subtract a fixed number of decibels, normalize so the loudest sample lands at a peak target you choose, or normalize so the integrated loudness matches a LUFS target. Peak and LUFS modes measure the material first and then apply exactly the right amount of gain to reach the destination, so you don't have to guess. A live transfer plot at the top of the panel shows what the current setting will do.

When to use it. Reach for fixed Gain to nudge a quiet take up or a hot one down. Use Peak to bring a file up to a safe ceiling just below clipping. Use LUFS to hit a delivery loudness for streaming, podcast, or broadcast. Use Gain for a fixed, peak, or quick loudness move; when you need to meet a broadcast or streaming LUFS spec with true-peak limiting, reach for **Loudness Control** instead.

Controls

- **Mode tabs (Gain / Peak / LUFS)** — choose how level is set: fixed gain, peak normalize, or LUFS target. The control below and the visualization change with the mode.
- **Gain** (*Gain mode*) — boost or cut the overall level by this many decibels (–48 to +48 dB, default 0). The plot draws the transfer line and flags where a boost would push samples past 0 dBFS.

- **Peak** (*Peak mode*) — scale so the loudest sample reaches this level (–24 to 0 dBFS, default –0.1). The plot shows the target line and the remaining headroom up to 0 dBFS.
- **Target** (*LUFS mode*) — adjust gain so integrated loudness matches this value (–31 to –6 LUFS, default –16), with a ± 1 LU tolerance band and ticks at the common targets.
- **–14 / –16 / –23** (*LUFS mode*) — one-tap buttons that set the Target to those common streaming and broadcast loudness values.

How to use it

1. Make a selection to process just that region, or leave nothing selected to process the whole file.
2. Pick a mode tab and set its value, or use a LUFS quick button.
3. Use **Preview** to audition, then press **Process** in the footer to commit — whichever mode is selected, Process applies it.
4. To queue these settings into a processing chain instead of applying now, use the footer **+** menu's **Add to Module Chain**.

Tips


- In Gain mode, watch the red clip flag — if a boost would cross 0 dBFS, trim it or switch to Peak mode for a safe ceiling.
- Peak normalizing to a level you've already passed will lower the file, not just raise it.
- LUFS normalizing only changes overall gain; it does not limit peaks, so very dynamic material may still need a peak ceiling afterward.

Presets

- **Normalize Peak –1 dBFS** — bring the loudest sample to a safe –1 dBFS ceiling.
- **Normalize Peak –0.3 dBFS** — push closer to full scale while staying below clipping.
- **+6 dB Boost** — lift a quiet take by a fixed 6 dB.
- **–6 dB Trim** — pull a hot file down by a fixed 6 dB.

DC Offset

Re-center the waveform on zero — subtract the measured offset, or block DC with a gentle zero-phase high-pass.



What it does. DC Offset removes the constant bias that shifts a waveform off the zero line — a common artifact of cheap converters, miswired gear, or synthesis. The panel measures the current offset of every channel and shows it as a read-only readout, so you can see the problem (and confirm it is gone) before and after processing. Two modes do the removal. **Static** subtracts each channel's measured mean over the processed range — exact, and it touches nothing but the constant component. **High-pass** runs a gentle first-order DC-blocking filter (cutoff 2–10 Hz) instead; because it is applied zero-phase (forward and backward), it removes DC *and* slow sub-sonic drift without any phase distortion or waveform tilt.

When to use it. Whenever the waveform sits visibly above or below the center line, when Statistics reports a DC offset, or before edits that assume a centered signal — fades that should land on silence, clicks at cut points, or maximizing headroom before normalizing (a DC-shifted file clips earlier on one side). Use Static for a plain constant offset; switch to High-pass when the baseline *wanders* (drifting offset, thermal or contact drift on old transfers) rather than sitting at one value.

Controls

- **Measured DC** (*read-out*) — the current mean sample value of each channel over the selection (or the whole file), with its dBFS equivalent. This is exactly what Static mode will subtract. For very large files that stream from disk the measurement happens during Process instead.
- **Static / High-pass** — choose the removal method: exact mean subtraction, or the DC-blocking filter.

- **Cutoff** (*High-pass mode*) — corner frequency of the first-order blocker (2–10 Hz, default 5 Hz). Content above roughly twice the cutoff passes essentially untouched (better than 0.1 dB at 100 Hz with the default).

How to use it


1. Make a selection to fix just that region, or leave nothing selected to process the whole file.
2. Check the **Measured DC** readout — it shows the per-channel offset the module found.
3. Pick **Static** (constant offset) or **High-pass** (drifting baseline), then **Preview** and **Process**.
4. Re-open the readout after processing — it should show *(none)*.

Tips

- Static subtraction is bit-transparent to the audio content: it only shifts the whole range down by the measured constant, so it is always the safest first choice.
- A file whose offset *changes* over time will still show residual DC after Static (the mean of a drifting bias is only its average) — that is the High-pass mode's job.
- DC removal is broadband by design: it always works on the full frequency range of the selected time span, even when a spectral (frequency-limited) selection is active.
- Remove DC before normalizing or limiting — the offset eats headroom on one side of the waveform.

Loudness Control

Hit an exact LUFS target while keeping true peaks under your ceiling.



What it does. Loudness Control measures the integrated loudness of your selection and applies a single, even gain move so the whole passage lands on the loudness target you choose. It then runs a true-peak limiter so the level move never pushes inter-sample peaks above your ceiling. The result meets the delivery spec you need — broadcast, streaming, podcast — without you guessing at gain.

When to use it. Reach for it when a master, episode, or stem needs to match a platform's loudness requirement (for example -14 LUFS for streaming or -23 LUFS for EBU R128). It is also the fastest way to bring a quiet recording up to a consistent, spec-compliant level. Use Loudness Control for broadcast and streaming LUFS delivery with true-peak limiting; for a plain fixed, peak, or quick loudness move without limiting, use the **Gain** module instead.

Controls

- **Loudness Target Scale** — a vertical dB ruler that shows, from your settings alone, where the module aims: a bright line marks the target, a translucent band brackets the $\pm$ tolerance window, and an amber rule near the top shows the true-peak ceiling with the headroom shaded. Faint ticks relate your target to common standards. It is purely a guide and never touches the audio.

- **Loudness standard...** menu — pick a preset spec (EBU R128, ATSC A/85, Streaming, Podcast, Netflix dialog, YouTube, Spotify loud, Sony games) to fill the target, ceiling, and tolerance in one move.
- **Target** — the integrated loudness the audio is gained to match. Lower is quieter, higher is louder (−31...−6 LUFS, default −16).
- **Tolerance** — how far from the target the measurement may sit before it is flagged as out of spec (0...3 LU, default 1).
- **Ceiling** — the maximum true-peak level the limiter will not exceed (−6...0 dBTP, default −1).
- **Limit true peak** — keeps inter-sample peaks under the ceiling after the loudness move (on by default).
- **Analyze** — measures current loudness and peaks without changing the audio.
- **Measurement panel** — after Analyze, shows the measured LUFS, the gain needed to reach target, the current and projected true peak, and a "within / outside tolerance" chip.

How to use it

1. Select the region (or whole file) you want to conform.
2. Pick a **Loudness standard**, or set **Target**, **Tolerance**, and **Ceiling** by hand.
3. Click **Analyze** to see where you stand and the gain that will be applied.
4. Confirm **Limit true peak** is on if you need a hard ceiling, then click **Apply Control**.
5. The panel re-measures automatically so you can verify the result lands within tolerance.

Tips


- Use **Analyze** first: the "After" reading tells you whether the limiter will engage before you commit.
- If the projected peak already sits under your ceiling, the limiter stays out of the way — you get a clean, transparent gain move.
- Loudness Control applies a single overall gain, so it preserves your dynamics; use it after, not instead of, dynamics processing.
- Tighten **Tolerance** (for example 0.5 LU) when a platform's QC is strict; loosen it for casual targets.

Presets

- **Streaming −14 LUFS** — conform to the common −14 LUFS streaming level with a −1 dBTP ceiling.
- **Podcast −16 LUFS** — podcast delivery at −16 LUFS with a −1.5 dBTP ceiling.
- **Broadcast EBU R128 −23** — broadcast spec at −23 LUFS, −1 dBTP, with tight 0.5 LU tolerance.
- **Club / Loud −9** — push to a hot −9 LUFS with a near-full −0.3 dBTP ceiling.

Trim Silence

Tighten the gaps — remove or shorten quiet passages automatically.



What it does. Trim Silence scans your audio and finds every passage that sits below the silence threshold for long enough to count as a gap. It then removes those gaps entirely, or shortens the long ones to a length you choose, splicing the kept audio back together with a short crossfade so the joins are clean. Silence at the very start and end of the audio is always removed in full; interior gaps follow your settings.

When to use it. Reach for it to tighten a podcast or interview after editing, to strip dead air from voice-over takes, or to remove long silent stretches from a recording before mastering. It works on the whole file or on just the selected range.

Controls

- **Threshold** — the level below which audio is treated as silence. Lower (more negative) values only catch truly quiet passages; higher values trim more aggressively and may cut into soft tails or room tone. (–100 to –12 dB, default –60 dB.)
- **Min** — the shortest gap that qualifies for trimming. Raise it to leave natural short pauses alone; lower it to catch brief gaps too. (10 to 2000 ms, default 250 ms.)
- **Padding** — how much audio to keep on each side of every kept region, so breaths and tails are not clipped at the joins. Higher values are safer but leave more silence. (0 to 500 ms, default 50 ms.)
- **Silence** — shorten long interior gaps to this length instead of deleting them outright. Set it to 0 to remove qualifying gaps completely. (0 to 2000 ms, default 0 ms.)

- **Crossfade** — the fade length applied around each splice point to avoid clicks. (0 to 100 ms, default 10 ms.)
- **Trim Silence** — applies the trim using the settings above.

How to use it

1. Select a range, or leave nothing selected to process the whole file.
2. Set **Threshold** and **Min** so the panel targets your dead air, not soft speech.
3. Set **Padding** to protect breaths, choose whether to remove gaps fully or shorten them with **Silence**, and pick a **Crossfade**.
4. Click **Trim Silence** to apply. Use **Add to Module Chain** to keep these settings as a step you can reuse.

Tips


- If words or breaths get clipped, raise **Padding** or lower **Threshold** before changing anything else.
- Use **Silence** when you want to keep a natural rhythm — shortening pauses sounds more human than deleting them.
- Start from a preset, then nudge **Threshold** to match your room tone.

Presets

- **Tight** — aggressive trimming for clean, close-mic'd material.
- **Podcast Dialogue** — keeps conversational pacing with generous padding.
- **Gentle** — conservative settings that only touch long, obvious gaps.

Harmonic Clarity

Reshape the balance of tone, air and attack — without changing the level.



What it does. Harmonic Clarity separates your selection into three kinds of energy — stable tonal partials, broadband noise/air, and percussive transients — and lets you rebalance them with five bipolar controls. You can clean up around a pitch or add air, sharpen or smear the partials, and bring out the attack or the sustain. Every move is loudness-matched, so each control changes the *character* of the sound, not its volume. FFT detects the dominant pitch of the selection automatically — there's no analyse button.

When to use it. Reach for it to add definition and air to a dull vocal or instrument, to smooth a harsh or metallic tone, to make a soft performance punchier, or to clean tonal-vs-noise content on material with no clear pitch.

Controls

A schematic "harmonic comb" at the top previews how your settings reshape a tone and its air floor. Below it, a **Selection** readout shows the detected **Pitch** plus meters for **Harmonic**, **Tonal**, **Noise / Air** and **Transient** content. Each control has a pill switch beside it (on by default); turning a control off zeroes it and skips that step, speeding up the render. All sliders span -2 to +2.

- **Harmonic** — left = inharmonic/metallic (spreads energy off the harmonic series); right = pure/harmonic (locks onto the detected overtones and clears clutter between partials). On atonal material it acts as a tonal-vs-noise cleaner. (default +0.6)
- **Focus** — left = smear partials into an ambient wash; right = sharpen spectral ridges for definition. (default 0)
- **Transients** — left = smoother, more sustain; right = bring out attacks (punch, pick, consonants). (default 0)
- **Consonance** — left = edgy/tense (boosts roughness of close partials); right = sweeter, more consonant. (default 0)
- **Beating** — left = steady the amplitude pulsing of close/detuned partials; right = enhance the beat for movement. (default 0)

How to use it

1. Select the region (or whole file) you want to treat.
2. Watch the Selection readout to see what the audio is made of.
3. Enable the controls you want and dial them in.
4. Press **Preview** to audition the reshaped sound.
5. **Process** to apply, or **Add to Module Chain** to stack it with other processes.

Tips


- Leave a control centred (or switch it off) when you don't need it — fewer active controls render faster.
- On noisy or atonal sources, lean on **Harmonic** as an air/clean control.
- Because everything is loudness-matched, compare settings by tone alone — level won't fool you.

Presets

- **Air & Definition** — adds air and sharpens partials with a touch of punch.
- **Smooth & Warm** — softens and warms a harsh or edgy tone.
- **Punchy** — accentuates attacks for a more percussive, present sound.
- **De-harsh** — tames roughness and steadies beating to calm harshness.

Time & Pitch

Stretch duration and shift pitch independently, with high-quality results.



What it does. Time & Pitch changes how long your audio plays and how high or low it sounds — and the two are completely independent. You can slow a clip down without dropping its pitch, raise it an octave without changing its length, or do both at once — and shift the *timbre* (the spectral envelope) separately from the pitch with the Formant control. It runs offline through the Signalsmith Stretch engine, replacing your selection (or the whole file when nothing is selected) with the stretched, shifted result in one undoable step.

When to use it. Tighten a voiceover to hit a target duration, drop a sample an octave for a bass part, nudge a take into tune, or create a subtle detuned doubling layer. It also handles the classic "fit this clip to that length" problem in post.

Controls

- **Time** — output duration as a percentage of the input (50–200%, default 100%). Above 100% lengthens (slows); below 100% shortens (speeds up); 100% leaves length unchanged.
- **Pitch** — coarse pitch shift in semitones (–12 to +12 st, default 0). Positive is higher, negative is lower.

- **Formant** — shifts the spectral envelope — the timbre or vowel colour — in semitones, independent of the pitch shift (–12 to +12 st, default 0 = unchanged). Counter-shift it to keep a pitched-up voice from turning chipmunk, or move it alone to change a voice's character at the same pitch.
- **Selection / File readout** — a live grid for the audio being processed: **Duration** (before → after in seconds), **Stretch** (100% → your Time setting) and **Tempo** (the Original BPM → the resulting BPM).
- **Cents** — fine pitch shift (–100 to +100 ct, default 0). 100 cents equals one semitone, so this trims pitch between the semitone steps.
- **Tonality Threshold** — frequencies below this limit keep their natural harmonic relationships for a more musical shift (0–16000 Hz, default 8000). Set it to 0 for a plain linear shift.
- **Original BPM** — the source's reference tempo (20–300 BPM, default 120). It only drives the resulting-tempo readout; it does not affect the processing.

How to use it

1. Select the region to process, or select nothing to treat the whole file.
2. Set **Time**, **Pitch**, and **Cents** to your target; watch the duration and tempo readout.
3. Adjust **Tonality Threshold** if a shift sounds unnatural, and **Formant** if a shifted voice sounds the wrong size.
4. Click **Process** to commit. There is no live preview because the length changes; undo if you want to try again.

Tips


- Process stays disabled until you change something away from the neutral identity (100% / 0 st / 0 ct / 0 st formant).
- Set **Original BPM** to the track's real tempo and the readout tells you the exact BPM the stretch will produce — handy for hitting a target tempo.
- Markers and regions are rescaled to follow the stretch automatically.
- For pure pitch moves, leave **Time** at 100%; for pure time-stretch, leave **Pitch** and **Cents** at 0.
- Because it changes the audio's duration, Time & Pitch cannot join the Module Chain or run in Batch.

Presets

- **Octave Up** — raise pitch a full octave, length unchanged.
- **Octave Down** — drop pitch a full octave, length unchanged.
- **Perfect Fifth Up** — shift up a perfect fifth (7 semitones).
- **Chorus Detune** — a gentle –8-cent detune for doubling and width.

Music Tailor

Fit a piece of music to an exact length with edits you can't hear.



What it does. Music Tailor re-cuts a piece of music to the length you need — shorter or longer — using musically hidden edits instead of a fade. It first analyzes the whole track for tempo, key and structure; then, for your target length, it proposes three edited versions. Shortening jumps forward at a matching point in the song; extending loops a compatible section. Every seam is beat-quantized and chord-aware — the chords are read on the song's own beats, with their quality and bass note, and the bar and the sections come out of that same reading — so a proposal can be a *whole section* ("repeat the chorus", "drop the second verse"), and a cut that falls on a drum hit can be a few milliseconds long instead of a long crossfade. Each version can be auditioned in full — or just at the transition — before you commit, and an optional slight time-stretch nudges the chosen version to land on the target exactly (confined to one calm passage where possible, so the rest of the result is untouched).

When to use it. Whenever a track has to hit a runtime: a 30-second ad cut of a 3-minute song, a podcast bed that must cover the intro, a film cue trimmed to the scene, or a loop-extended piece that has to carry a longer edit. It replaces the manual find-a-bar / razor / crossfade ritual with proposals you simply audition.

Controls

- **Status chip** — walks you through the state: "Analyze the track to begin", "Analyzing...", "Analyzed — set a length and find versions" (tagged READY), "Finding versions...", "Rendering..."
- **Workflow strip** — the three steps ticked off as you go: **Analyze** → **Set length** → **Choose version**.

- **Analyze / Re-analyze** — scans the whole track for tempo, key and structure; a **Cancel** button appears while it runs. After the audio changes, the proposals go stale and the panel asks you to analyze again.
- **Pro (lyrics aware)** — off by default, the module's one option. It adds two slow passes to Analyze: it separates the voice from the backing so that seams are steered away from sung phrases, and it reads the chords with the AI model. Expect a few minutes on a full song — and it reuses work another module has already done: if Choir Maker or Lyrics & Chords separated this file, Pro is free, and stems Pro pays for are handed back to them. Leave it off for a quick pass; turn it on for anything with a lead vocal.
- **Analysis pills** — the detected BPM (a single figure, or a range like "112–128 BPM" when the tempo genuinely moves), the detected key, and a tempo-stability verdict: *steady*, *varying tempo* or *loose tempo*.
- **Target length** — a scrubbable slider in seconds paired with an exact **m:ss** text field (accepts `3:05`, `3:05.5` or plain seconds), with a readout of the source length beside it. It opens pre-filled with the source length.
- **Slight time-stretch (0.9×–1.1×)** — nudge each version toward the target with a gentle stretch. Quality comes first: versions beyond the $\pm 10\%$ range keep their natural length instead of forcing a rough cut.
- **Mark edits in result** — document the edits in the rendered file: markers at each cut, and loop passes become named regions.
- **Protect selection** — on by default, and it costs nothing until you select something. With a passage selected, no seam may land inside it and no cut or loop may cover it: the selection survives the edit intact and exactly once. Markers named `keep...` and `cut...` are honoured too, as gentle preferences rather than rules — a `cut here` marker attracts seams, a `keep` marker pushes them away.
- **Find 3 versions** — proposes three different edits that fit the target (enabled once the analysis is fresh).
- **VERSIONS list** — one row per proposal, titled by its plan and final length: "Shorten — 2:30.00", "Shorten · 2 cuts — ..." or "Extend ×3 — ...". A row whose edit is a whole detected section says so — "Extend ×2 (repeats a section)", "Shorten (removes a section)" — which is usually the version to try first. Each row shows its seams ("jump 1:42.30 → 2:05.10" for a cut, "loop 2:05.10 ↔ 1:42.30" for an extension) and a quality line — transition match % (the weakest final simulated crossfade; versions below 95% are labelled "best available"), arc % (how well the song's shape survives) and, when stretching, the stretch %. Per row: a **locate** button that selects the edited span on the waveform and scrolls it into view (click it again on a multi-cut version to step through its cuts one at a time, then the whole span it touches), a **play** button that auditions the whole version, a looping **transition** audition per seam (numbered when there are several), and **Apply**, which renders that version as one undoable edit. Changing the target invalidates old versions, so Apply can never reuse a plan made for another duration.

How to use it

1. Open the track and click **Analyze** — check the pills say *steady* for the most reliable results.
2. Set the **Target length** with the slider or type it exactly (e.g. `2:30`).


3. Decide on **Pro (lyrics aware)** before analyzing if the track has a vocal, then on **Slight time-stretch** (exact length), **Mark edits in result** (visible cut points) and **Protect selection** (select the passage that must survive first).
4. Click **Find 3 versions** and audition each proposal — the looping transition button is the fastest way to judge a seam.
5. Click **Apply** on the winner. The render is one undoable edit, so you can ⌘Z (Ctrl+Z on Windows) and try another version.

Tips

- Music Tailor works best on music with a steady tempo; on *varying* or *loose tempo* material the seams are harder to hide, so audition transitions carefully.
- The transition audition loops until you stop it — let it cycle a few passes; a good seam disappears on repetition.
- Without time-stretch, versions land *near* the target (the list shows each version's offset, e.g. "+1.2s vs target"); with it, they land exactly on it.
- Use the locate button before applying to see exactly which span gets removed or looped.
- On a song with a lead vocal, turn **Pro (lyrics aware)** on before you Analyze: the seam a listener notices first is one that lands mid-word, and the separated voice is what lets Music Tailor weigh that against every other candidate. It is a strong preference, not a veto: far fewer seams land in the vocal, and auditioning each transition is still the last word.
- Music Tailor is analysis-only: there is no Preview/Process footer — **Apply** on a row is the render — and it is not available in the Module Chain or Batch.

Loop Creator

Turn any passage into a seamless, bar-length loop — found in the audio or grown from it.



What it does. Loop Creator makes a clean, repeating loop of an exact musical length — one, two, four or eight bars at the track's tempo. It works two ways. **Find** scans the recording (or your selection) for a bar-length section that already loops back on itself smoothly and offers the best candidates, ranked. **Generate** synthesizes a brand-new loop with on-device Stable Audio 3, conditioned on the audio surrounding the start point so it matches the feel of the track — no text prompt needed. Either way, the result is written back as a seamless loop, in place or into a new document.

When to use it. Build a backing loop, an ambience bed, or a rhythmic underlay from existing material; extract a tidy two-bar groove to repeat under a voiceover; or grow a fresh loop that sits naturally in the same key and texture as the source. It replaces the manual hunt-for-a-zero-crossing / trim-to-the-bar / crossfade ritual with ranked proposals you can audition on repeat.

Controls

- **Status chip** — tracks the state: "Analyze to detect the tempo", "Analyzing tempo...", "Analyzed — set bars and create a loop" (tagged READY), then "Finding loops..." / "Generating loop..."
- **Mode picker** — **Find** (scan the audio for a loop) or **Generate** (synthesize one). The spectrogram brackets the surrounding audio a Generate loop conditions on.

- **Analyze / Re-analyze** — detects the tempo (BPM) and beat grid; a **Cancel** appears while it runs. The detected figure shows beside the button. After the audio changes the analysis goes stale and the panel asks you to analyze again.
- **BPM** — the working tempo, seeded from the detection but fully editable — type or scrub it if the detector locked onto a half- or double-time reading.
- **Bars** — the loop length in bars: $\frac{1}{2} \cdot 1 \cdot 2 \cdot 4 \cdot 8$. The line beneath spells out the resulting length, e.g. "4 bars · 8.00 s at 120 BPM".

Find mode

- **Find loops** — scans the whole file, or just the selection ("in selection" / "whole file" is shown), for seamless loops of the chosen bar length.
- **LOOPS list** — up to three ranked proposals. Each row shows its **match %** and a detail line — its start time, a **seam** score (how cleanly the end meets the start) and a **rhyme** score (how self-similar the content is). Per row: a **scope** button that selects the loop on the waveform, an auto-looping **preview** (repeats until you stop it), and **Create Loop**.

Generate mode (behind the Stable Audio 3 download gate)

- **Surrounding audio** — how many seconds around the start point condition the synthesis (0–15 s, default 4). The spectrogram brackets exactly this span.
- **Generate loop** — synthesizes a bar-length loop grown from the surrounding audio (no text prompt). A **Cancel** appears while it runs.
- **Generated loop row** — the result, with its start time, an auto-looping **preview**, and **Create Loop**.

Both modes

- **Open in a new document** — when on, Create Loop opens the loop as a fresh tab; when off, the loop replaces the current selection (or is placed at the cursor).

How to use it

1. Click **Analyze** to detect the tempo, and correct the **BPM** if it locked onto half- or double-time.
2. Pick the **Bars** length you want the loop to be.
3. **Find:** click **Find loops**, then preview the ranked proposals — the auto-looping button is the quickest way to judge a seam. **Generate:** set **Surrounding audio**, click **Generate loop**, and preview the result.
4. Decide placement with **Open in a new document**, then click **Create Loop**.


Tips

- Find looks for material that already loops, so it shines on steady, groove-based passages; if nothing scores well, try a different bar length or a tighter selection.
- If the loops feel like they run at the wrong speed, the tempo detection probably caught a half/ double-time beat — halve or double the **BPM** and find again.
- The preview repeats until you stop it — let it cycle a few passes; a good seam disappears on repetition.

- Generate needs the Stable Audio 3 model installed, its model terms accepted, and an active license; Find is plain editing and works without any of these.
- Loop Creator is analysis-only: there is no Preview/Process footer — **Create Loop** is the commit — and it is not available in the Module Chain or Batch.

Dither

Reduce bit depth cleanly, trading harsh quantization distortion for a faint, benign noise floor.



What it does. Dither lowers the bit depth of your audio to a target word length. Whenever you reduce bits, the values have to be rounded to a coarser grid, and that rounding can add gritty, level-dependent distortion to quiet passages and fade tails. Dither sidesteps that by adding a tiny amount of triangular (TPDF) noise before rounding, which turns the distortion into a steady, natural-sounding hiss your ear ignores. An optional noise-shaping stage moves that residual noise up into higher frequencies, where your hearing is less sensitive, so it disturbs the music even less.

When to use it. Use it as the final step before delivering a 16-bit (or other reduced-depth) file from a 24-bit or floating-point project, especially on material with quiet fades, reverb tails, or solo instruments. If you are exporting straight to 16/24-bit PCM, the export dialog can apply the same dither for you; the module is for dithering in place, auditioning the result, or placing dither inside a chain.

Controls

- **New bit depth** — The target word length to requantize to. Choose from **24**, **20**, **16**, **12**, or **8-bit** (default **16-bit**). Fewer bits means a coarser grid and a more audible noise floor.
- **Dither amount** — How much triangular noise is mixed in before rounding. **None** truncates with no dither (only sensible alongside noise shaping); **Normal** adds a standard one-step amount; **High** adds twice that for extra distortion masking at the cost of a slightly louder hiss (default **Normal**).
- **Noise shaping** — When on, pushes the residual noise toward high frequencies so it is harder to hear (default **on**).

How to use it


1. Select the region to process, or leave nothing selected to treat the whole file.
2. Set **New bit depth** to your delivery target.
3. Pick a **Dither amount** (start with Normal) and leave **Noise shaping** on.
4. Click **Process** to apply. Use **Add to Module Chain** to keep dither as the last stage of a saved chain.

Tips

- Always make dither the final process — any gain, EQ, or limiting after it re-introduces rounding error.
- Dither once. Repeatedly dithering the same audio only stacks noise.
- None + Noise shaping is the cleanest choice when you want minimal added energy and can accept pure shaped truncation.
- Going to 8 or 12-bit is inherently noisy; High amount with shaping masks distortion best at those depths.

Plugin Host

Run your installed plug-in effects — Audio Units on macOS, VST3 on Windows — inside FFT, baked straight into the file.



What it does. Plugin Host lets you process audio through the third-party effects you already have installed on your computer. It scans your system for plug-ins and lists them in a menu; pick one, tune it in its editor, and FFT renders the selection through it offline, time-aligned, and writes the result back like any other module. The plug-in format is the platform's own: on macOS FFT hosts **Audio Unit** effects (VST3 bundles are listed but cannot be hosted there — the panel says so if you pick one); on Windows FFT hosts **VST3** effects through its own crash-isolated scanner, and any plug-in the scan had to skip (an x64-only plug-in in an ARM64 host, for example) is named with the reason rather than silently dropped.

When to use it. Reach for it when the effect you want lives in a plug-in you own rather than in one of FFT's built-in modules — a favourite compressor, a specialist de-noiser, a saturation or EQ you trust. Add it to a Module Chain when you want to reuse the same plug-in treatment; direct Batch catalog entries exclude Plugin Host, but Batch can run supported chain steps when the plug-in is available.

Controls

- **Plugin** — the dropdown of discovered effects (on macOS, Audio Units appear as "AU — ..." and VST3 bundles as "VST3 — ..."; on Windows the list holds the VST3 plug-ins the scan found); choose **None** to disable the module. Switching plug-ins clears any custom settings from the previous one. (Default: None.)

- **Scan installed plug-ins** (the circular refresh button on macOS; **Rescan** on Windows) — rescans your system so newly installed effects show up without reopening the panel.
- **Open Plug-in...** — shows the selected plug-in's editor so you can tune it (inline in the module panel on macOS; in the plug-in's own window on Windows, after you press **Load**). Every change is saved automatically and travels with presets, module-chain steps, and renders. If the plug-in has no interface of its own, FFT shows a generated list of its parameters as sliders. ("Custom settings saved" confirms your tweaks are stored.)
- **Reveal** (macOS) — shows the plug-in file in Finder.
- **Mix** — wet/dry blend, from fully dry to fully processed (0–1, default 1).
- **Output** — output gain trim applied after processing (–24 to +24 dB, default 0).

How to use it


1. Make a selection (or leave it empty to process the whole file).
2. Pick an effect from **Plugin**; press the refresh button if it isn't listed yet.
3. Click **Open Plug-in...** and set it up in the inline editor.
4. Set **Mix** and **Output** to taste.
5. Click **Process** to commit. Use **Add to Module Chain** to stack it with other modules.

Tips

- Only *effect* plug-ins are hostable — standard Audio Unit effects on macOS, VST3 effects on Windows. On macOS, selecting a VST3 entry makes the panel tell you VST3 hosting isn't supported there.
- A misbehaving plug-in can never corrupt your audio: any failure leaves the file untouched and shows a message, and at **Mix** 0 the original passes through untouched.
- Your plug-in tuning is captured even after you close its editor, so presets and supported module-chain runs reproduce it exactly.

Match

Learn a recording you like, then steer your audio toward it — its tone alone, or the whole master.



What it does. Match makes your audio resemble a reference recording, in one of two ways. **EQ curve** mode studies the reference's overall tonal balance — the long-term spread of bass, mids and treble — and applies a fixed, broad EQ curve that nudges your audio toward that same balance. It matches *tone*, not loudness: the correction is gentle, smoothed across 48 bands, and phase-preserving. **Full master** mode learns the reference's broader sonic profile — overall loudness, low and high tonal balance, and stereo width — and applies a matched gain, a low shelf around 180 Hz, a high shelf around 4.2 kHz, a stereo width adjustment and a peak-safe output ceiling: a light "make it sit like this master" pass, not a note-for-note tonal copy.

When to use it. EQ curve mode when one clip should sit alongside another — matching a re-recorded line into a dialogue scene, making a dull take sound like a brighter one, or applying a "house tone" across many clips. Full master mode when you have a finished reference master and want your mix to land in the same ballpark for loudness, brightness, low end and width — matching a single across an album, or aligning a new master to a client's approved reference.

Controls

- **EQ curve | Full master** — the mode picker at the top. Each mode keeps its own learned reference, and the rest of the panel follows the choice.
- **Workflow steps** — a small guide (Learn a reference → Adjust amount → Render) that ticks off as you go.
- **Status chip** — the name of the learned reference, or "No reference learned".

- **Learn Selection** (*EQ curve*) — uses your current selection (or the whole file) as the tonal reference.
- **Learn File...** (*EQ curve*) — opens a picker to use a separate audio file as the reference.
- **Learn Reference...** (*Full master*) — opens a file picker; the full-master reference always comes from a file.
- **Clear** — appears once a reference exists; discards the current mode's reference so you can start over.
- **Amount** — how strongly the reference is applied. In EQ curve mode 0–200% (default 100%); 100% matches the reference's tone closely and above that over-corrects, with boosts capped at +12 dB and cuts at –18 dB. In Full master mode 0–100% (default 100%).
- **Ceiling** (*Full master*) — the peak-safe output ceiling applied after matching (–12 to 0 dB, default –1.0 dB).
- **Correction curve** (*EQ curve*) — a bar strip showing the per-band move the current settings would make: orange bars for boosts, cyan for cuts. It updates live with the selection and Amount.
- **Render applies** (*Full master*) — a live summary grid of exactly what Process will do: the gain (an RMS match), the low shelf (180 Hz), the high shelf (4.2 kHz), the stereo width factor, and the ceiling.

How to use it

1. Pick the mode: **EQ curve** for tone only, **Full master** for level, tone, width and ceiling.
2. Learn a reference — **Learn Selection** / **Learn File...** in EQ curve mode, **Learn Reference...** in Full master mode.
3. Set **Amount** (and **Ceiling** in Full master mode), watching the correction curve or the "Render applies" grid.
4. Click **Process** to apply the match as one undoable edit (EQ curve or Full master mode).

Tips


- Learn the reference from a clean, representative passage — a quiet or atypical section will skew the match.
- Start at 100% and ease the Amount down if the result feels heavy-handed; small moves sound the most natural.
- EQ curve mode matches tone and ignores level — pair it with a gain or normalize step if you also need loudness. Full master mode includes the level: its gain is set so peaks stay under your Ceiling, preserving dynamics without limiter pumping.
- The "Render applies" grid is your safety check — very large values usually mean a very different reference. Stereo width only changes stereo material; mono files are matched for level and tone only.
- Only EQ curve mode joins the Module Chain and Batch. Full master mode is render-only and the chain declines it ("Not chainable in Full Master mode") — switch to EQ curve to chain a match.

AI Modules

On-device machine-learning tools. Models must be bundled with the app or installed locally first, then run on your computer's GPU (Apple Silicon GPU / Neural Engine on macOS, your GPU or CPU on Windows) with no cloud upload.

Stem Remix

Split a finished mix into stems, then rebalance or export them.



What it does. Stem Remix uses on-device AI to separate a stereo mix into four stems — Drums, Bass, Other, and Vocals — and gives each its own level fader. Raise or lower any stem and FFT re-renders the mix to your new balance, all without leaving the editor. You can also save the four separated stems out as individual audio files. Separation runs locally on the GPU; nothing is uploaded.

When to use it. Reach for it when you only have a final mixdown but need control over its parts: pull the vocal up in a rough mix, tame loud drums, lift a buried bass line, or duck everything but the voice. It's equally useful for grabbing clean stems to remix, sample, or reuse elsewhere.

Controls

- **Status row** — shows the current state ("4-stem separation", "Separating into 4 stems", or "Stems ready") and confirms the engine is running on-device (the tag reads "Core ML GPU" on macOS and "On-device" on Windows). If the model isn't installed yet, an install hint appears with the location the converted weights are expected in.
- **Drums / Bass / Other / Vocals faders** — one level fader per stem (range -24 dB to +12 dB, default 0 dB each). Drag up to push a stem forward, down to pull it back; drag fully to the bottom and the readout shows **mute** to drop that stem entirely. Each fader's value is shown live in dB.
- **Separate & Apply / Apply Rebalance** — the main action button. On the first run it reads **Separate & Apply**: it separates once, then renders the rebalanced mix as one undoable edit. Once stems already exist it reads **Apply Rebalance** and re-renders instantly with the current fader levels.
- **Reset** — returns all four faders to 0 dB, i.e. the unchanged original mix.
- **Separate Only / Re-separate** — splits the audio into the four stems without touching the mix, so stems are ready to export. Labels switch to "Re-separate" once a separation already exists.

- **Export Stems...** — saves each separated stem to its own audio file in a folder you choose. Available only after a separation exists.
- **Cancel** — appears while separating, to stop the running job.

How to use it


1. Select a region, or leave nothing selected to process the whole file.
2. Set the four stem faders to the balance you want.
3. Click **Separate & Apply** (first run) — it separates once, then renders the rebalanced mix as one undoable edit.
4. Tweak the faders again and press **Apply Rebalance**: with stems already ready, it re-renders instantly.
5. To keep the parts as files instead, use **Separate Only** then **Export Stems....**

Tips

- Separation happens once. After it's done, fader changes re-render instantly, so audition several balances freely.
- Editing the audio or changing the selection invalidates the stems — you'll need to separate again.
- The main action button stays disabled until at least one fader differs from 0 dB, since a flat mix would change nothing — the hint line under the buttons tells you so while it's greyed out.
- Muting a stem (fader fully down) is the quickest way to isolate or remove a single element.

Voice Isolate

Split a track into vocal and instrumental — keep either, blend them, or export both.



What it does. Voice Isolate uses an on-device Mel-Band RoFormer model to separate your audio into two stems — the vocal and everything else. Pick which stem to keep and FFT renders it back over the selection (or the whole file) as one undoable edit, with a dry/wet Amount so you can blend the isolate against the original instead of committing to a hard split. Both stems can also be saved out as files. Separation runs locally on the GPU; nothing is uploaded.

When to use it. Pull a clean vocal out of a finished mix for a remix or transcription, make an instant instrumental/karaoke version, duck the backing to feature the voice, or use a partial Amount to simply push the vocal forward without fully soloing it.

Controls

- **Status row** — shows the current state ("Vocal / instrumental separation", "Isolating vocal / instrumental", or "Stems ready") and confirms the engine runs on-device (the tag reads "MLX GPU" on macOS and "On-device" on Windows). If the selected engine's model isn't installed yet, a download manager appears in its place.
- **Engine** — *Fast* (33 MB, the default) or *HQ* (228 MB, higher quality). Each variant is a separate model that downloads on demand the first time you select it.
- **Keep** — *Vocal* or *Instrumental*: which stem **Process** renders back into the file. Export always writes both stems regardless.
- **Amount** — dry/wet blend of the kept stem over the original (0–100 %, default 100 %). At 100 % you get the full isolate; lower values mix the original back in.
- **Process / Cancel** — separates into vocal / instrumental if needed, then renders the kept stem over the range as a single undoable edit. While the model runs, the button becomes **Cancel**. Once stems exist, Process re-renders instantly.

- **Export stems** — choose a folder, then FFT separates (if needed) and saves *<filename>-Vocals.wav* and *<filename>-Instrumental.wav* (32-bit float WAV).
- **Model credits & license** — opens the model's attribution and license details.

How to use it


1. Select a region, or leave nothing selected to process the whole file.
2. Pick an **Engine** — Fast for speed, HQ for the cleanest split — and let it download if prompted.
3. Choose what to **Keep** and set the **Amount**.
4. Click **Process**. The first run separates once; after that, changing Keep or Amount and pressing Process again re-renders instantly (undoable).
5. To save the parts as files instead, click **Export stems** and pick a folder.

Tips

- Separation happens once and is cached — audition Vocal vs Instrumental, or several Amounts, without re-running the model.
- Editing the audio or changing the selection invalidates the stems; the next Process or Export separates again.
- An Amount below 100 % is often more natural than a hard isolate — leaving a little of the original masks separation artifacts.
- Export doesn't touch the document at all, so it's the safe way to grab both stems for use elsewhere.

Dialogue Rebalance

Separate speech, music and effects — then remix their levels or export the stems.



What it does. Dialogue Rebalance uses the on-device Cocktail-Fork MRX model to separate a soundtrack-style mix into three stems — Speech, Music and SFX — and gives each its own level fader. Set the faders to the balance you want and FFT separates once, then renders the remixed result as a single undoable edit. The three stems can also be saved out as individual files. Everything runs locally on the GPU; nothing is uploaded.

When to use it. It's built for post-production dialogue problems: music or effects burying the dialogue in a video mixdown, a podcast bed that's too loud under the voices, or the reverse — pulling the dialogue down to reuse the music and atmosphere. It's also the quick way to extract a dialogue-only, music-only or effects-only stem from a finished mix.

Controls

- **Status row** — shows the current state ("Speech / music / sfx separation", "Separating speech / music / sfx", or "Stems ready") and confirms the engine runs on-device (the tag reads "MLX GPU" on macOS and "On-device" on Windows). If the model isn't installed yet, a download manager appears in its place.
- **Speech / Music / SFX faders** — one level fader per stem (−24 dB to +12 dB, default 0 dB each). Drag up to push a stem forward, down to pull it back; drag fully to the bottom and the readout shows **mute** to drop that stem entirely.
- **Process** — separates into speech / music / sfx if needed, then renders the mix with the current fader levels (undoable). It stays disabled while all three faders sit at 0 dB, since that would change nothing. Once stems exist, later fader changes re-render instantly.
- **Cancel** — appears while separating, to stop the running job.

- **Export stems** — choose a folder, then FFT separates (if needed) and saves each stem as `<filename>-speech.wav`, `<filename>-music.wav` and `<filename>-sfx.wav` (32-bit float WAV).
- **Credits & License** — opens the model's attribution and license details.

How to use it


1. Select a region, or leave nothing selected to process the whole file.
2. Set the **Speech / Music / SFX** faders to the balance you want — for clearer dialogue, try pulling Music and SFX down a few dB rather than boosting Speech.
3. Click **Process**. The first run separates once, then renders the rebalanced mix as one undoable edit.
4. Tweak the faders and press **Process** again — with stems ready, it re-renders instantly.
5. To keep the parts as files instead, click **Export stems** and pick a folder.

Tips

- Separation happens once and is cached, so audition several balances freely.
- Editing the audio or changing the selection invalidates the stems — the next Process or Export separates again.
- Muting a stem (fader fully down) is the quickest way to strip music or effects out of a dialogue recording entirely.
- Gentle moves keep it transparent: big boosts can expose separation artifacts, while cuts of a few dB usually pass unnoticed.

AI Speech De-noise

Deep-learning speech denoise running on-device.



What it does. AI Speech De-noise uses a deep-learning model, running entirely on your computer, to strip background noise out of recorded speech. Instead of profiling the noise the way a classic spectral de-noiser does, it has learned what a clean human voice sounds like and removes everything that isn't it — hiss, hum, room rumble, fans, traffic, keyboard chatter and more — in one pass. It works at full-band 48 kHz; files at other sample rates are converted in and back out so length is preserved exactly.

When to use it. Reach for it on dialogue, podcasts, interviews, voiceover and any single-voice recording that needs fast, hands-off cleanup. It is trained for speech only — for music or broadband material, use Spectral De-noise instead.

Controls

- **Engine** — picks the model. *Standard* is fast, suppression-style cleanup and is always available. *Studio* is a stronger full-band speech-restoration engine (MossFormer2-SE): rather than just suppressing noise it re-estimates the clean voice, which digs deeper into heavy noise. It appears

only when its model is installed; if the model is missing, the module falls back to Standard. (default: Standard)

- **Strength** (*Studio only*) — wet/dry blend of the restored speech against the original. Higher leans fully on the restored signal; lower keeps more of the original. (0–100%, default 100%)
- **Limit noise reduction** (*Standard only*) — caps how aggressively noise is pulled down, keeping borderline material natural rather than over-scrubbed. (on by default)
- **Max Reduction** (*Standard only, shown when the limit is on*) — the ceiling on how much noise can be removed. Lower is gentler and more transparent; higher allows deeper cleanup. (3–40 dB, default 12 dB)

How to use it

1. Select the speech region (or leave nothing selected to process the whole file).
2. Choose **Standard** for speed, or **Studio** for the deepest restoration if installed.
3. Set **Strength**, or enable **Limit noise reduction** and dial **Max Reduction**.
4. **Preview** to hear the result.
5. **Process** to commit, or **Add to Module Chain** to stack it with other modules.

Tips

- If voice sounds thin or "underwater," lower Max Reduction or pull Studio Strength back from 100%.
- Preview and A/B against the source often — if the cleanup starts eating breaths or words, ease off.
- Studio is heavier; Standard is more than fast enough for long files.

Spectral Recovery

Rebuild the missing highs in muffled, phone, or low-bitrate audio.



What it does. Spectral Recovery detects how much real high-frequency content your recording actually has, then re-synthesizes everything above that edge using an on-device algorithm. Band-limited material — phone calls, low-bitrate streams, dull or muffled captures — gets its top end back, so speech and music sound open and full-band again. If the audio is already full-band, the module short-circuits: rendering it changes nothing.

When to use it. Reach for it on conference-call recordings, voice memos, archival or telephone audio, and any source that sounds boxed-in or cut off at the top. It is at its best restoring clarity to speech, but also helps thin music and sound effects.

Controls

- **Status chip** — shows which algorithm is active and whether its model is installed. If a model is missing, a download manager appears in the panel to fetch it (Harmonic needs no model at all).

- **Input spectrum** — the selection's spectrum with the detected bandwidth **edge** marked by a dashed line; the band above the edge is shaded and labeled *rebuilt* — the band the render rebuilds. Full-band audio shows no shading, just the "full-band · passes through" readout.
- **Algorithm** — one menu with the five regeneration algorithms, ordered fastest to slowest:
 - *Fast (LavaSR)* — LavaSR v2, a lightweight speech bandwidth-extension net; by far the fastest model-backed option (~80× realtime).
 - *Harmonic* — model-free harmonic reconstruction (~50× realtime). Pure DSP: always available, nothing to download.
 - *Neural Speech (AERO)* — the AERO voice super-resolution net (~1.4× realtime).
 - *Neural Music (AERO)* — the AERO music super-resolution net (~0.5× realtime).
 - *Studio (MossFormer)* — MossFormer2-SR full-band regeneration (~0.4× realtime), the default. When the net can't contribute usable highs on a given render, Studio falls back to the model-free Harmonic rebuild instead of reporting a success that added nothing. Model-backed algorithms install on demand when first selected.
- **Amount** — the level of the re-synthesized highs (0–3, default 1). 1 is the model's natural output; values above 1 boost the rebuilt band (the models are conservative, so a lift often helps); 0 turns the regeneration off.

How to use it

1. Make a selection, or leave none to process the whole file — the spectrum display shows the detected edge and shades the band that will be rebuilt.
2. Pick the **Algorithm** — Harmonic or Fast when speed matters, Neural Speech for voice, Neural Music for musical material, Studio as the heavyweight all-rounder — and let it download if prompted.
3. Set **Amount**: start at 1, then raise it if the restored top end is too shy.
4. Click **Preview** to audition the rebuilt top end.
5. Click **Process** to commit, or **Add to Module Chain** to stack it with other modules.

Tips

- If Process does nothing, the source is already full-band — that's expected (the spectrum display will show no shaded region).
- Start at Amount 1 and adjust by ear: push above 1 for a brighter, more assertive rebuild, or back toward 0 if the new highs sound harsh or noisy.
- The same **Algorithm** menu appears in the **Resample** module, where it can rebuild the missing octaves after upsampling.

Voice Enhance

One-knob studio voice cleanup: AI denoise, de-reverb, tone, leveling and loudness in a single chain.



What it does. Voice Enhance runs your recording through a complete studio voice chain in one pass: AI speech denoise, de-reverb, automatic de-hum, de-plosive, de-ess, broadcast tone EQ, a speech leveler and a loudness target with a true-peak-safe limiter. A single master **Strength** knob scales every enabled stage, so you get a clean, consistent, podcast-ready voice from one control. Each stage stays individually toggleable and tunable when you want to fine-tune. It is tuned for spoken voice — podcasts, dialogue and voiceover.

When to use it. Reach for it on a raw spoken-word take that needs to sound finished fast: a noisy room recording, an echoey interview, a muffled phone or laptop-mic clip, or any dialogue that needs leveling and a broadcast loudness target.

Controls

- **Strength** — master macro that scales every enabled stage at once. Higher cleans more aggressively; at 0% the chain is a true pass-through (the loudness stage is also skipped). (0–100%, default 80%)
- **AI De-noise** — toggle plus amount. Removes background noise; at 100% the reduction is unlimited. (amount 0–100%, default 85%, on)
- **Engine** — appears only when AI De-noise is enabled. *Standard* is the bundled fast engine; *Studio* uses the strongest installed full-band restoration model and falls back to Standard until that model is installed (the tab reads "Studio (not installed)" in that case).
- **Regenerate** — re-synthesizes the frequency band above a muffled or phone recording's detected bandwidth; full-band audio passes through untouched. Requires the bandwidth-regeneration model. (amount 0–100%, default 25%, on)
- **De-reverb** — attenuates room reverb around the voice. (amount 0–100%, default 60%, on)
- **De-ess** — tames harsh sibilance. (amount 0–100%, default 50%, on)
- **Tone** — broadcast voice EQ: rumble cut, warmth, mud cut, presence and air. (amount 0–100%, default 60%, on)
- **Leveler** — rides speech toward a constant level and holds in pauses so noise never pumps up. (amount 0–100%, default 60%, on)
- **De-hum (auto)** — automatic hum/buzz removal. (toggle, on)
- **De-plosive** — softens "p" and "b" pops. (toggle, on)
- **Loudness target** — normalizes the result to a target with a true-peak-safe limiter. (toggle, on). When on, set **Target** (–24 to –10 LUFS, default –16) and **Ceiling** (–3 to –0.1 dB, default –1).

How to use it


1. Select the region (or all) you want to clean.
2. Leave the defaults, or adjust **Strength** to taste.
3. Toggle individual stages on/off and fine-tune their amounts if needed.
4. Set your **Loudness target** (–16 LUFS is the podcast standard).
5. Click **Process** to commit. Use **Add to Module Chain** to stack it with other modules.

Tips

- Preview and A/B against the source to confirm the chain is cleaning the voice, not stripping it — pull **Strength** down if it starts sounding thin.
- Studio engine and Regenerate need their models installed; until then those stages fall back or stay inactive.
- Pull **Strength** down for a natural, lightly-cleaned sound; push it up for heavy denoise on rough source.
- –16 LUFS suits podcasts; lower the Target for quieter, more dynamic delivery.

Transcribe

Turn speech into editable text and subtitles, entirely on your computer.



What it does. Transcribe listens to the spoken word in your audio and writes it out as time-stamped text, running completely on-device with GPU acceleration. The result is a clickable transcript: each segment carries a timestamp, and you can copy the full text, drop a timeline marker at every segment, or export an SRT subtitle file. If you have a selection active, only that region is transcribed; otherwise the whole file is.

When to use it. Reach for it to caption a podcast or video, to log an interview, to scrub long dialogue takes by reading instead of listening, or to navigate a recording by jumping straight to the line you need.

Controls

- **Model** — pick the speech engine. *English* is English-only and fast; *Multilingual* handles English plus other languages. Engines must be bundled with the app or installed locally before they can run; if one is missing, the panel shows its status and lets you check again after installation.
- **Language** — (*Multilingual model only*) tells the engine the spoken language. Leave it on **Auto-detect** (default) to sample the opening seconds, or pick from the list (English, Italiano, Español, Français, Deutsch, Português, Nederlands, 日本語, 中文, 한국어, Русский) when you already know it.
- **Translate to English** — (*Multilingual model only, off by default*) renders the transcript in English while transcribing, regardless of the spoken language.
- **Add markers for segments** — (*off by default*) drops a timeline marker at the start of each segment as the transcript is produced.
- **Transcribe / Cancel** — starts the run (the button reads "Transcribing" while it works) or stops it.
- **Transcript list** — the finished segments, each with a timestamp. Click any line to select that audio and move the playhead to it; the line under the playhead highlights during playback.

- **Create Markers / Copy Text / Export SRT...** — add per-segment markers after the fact, copy the whole transcript to the clipboard, or save a subtitle file.

How to use it

1. (Optional) Select the passage you want; leave nothing selected to transcribe the whole file.
2. Choose a **Model** and, for Multilingual, set the **Language** (or Auto-detect).
3. Toggle **Translate to English** and **Add markers** if you want them.
4. Click **Transcribe** and wait for the progress bar.
5. Click lines to navigate, then **Copy Text**, **Export SRT...**, or **Create Markers** as needed.

Tips

- Tighten a noisy or music-heavy mix with a denoise module first — cleaner speech transcribes more accurately.
- Use **Auto-detect** when a clip mixes languages or you're unsure; pin a specific language for the most reliable result.
- The transcript stays valid as long as the audio content is unchanged, so you can keep clicking and exporting without re-running.
- Exported SRT carries the same timestamps you see in the list, ready to drop onto a video timeline.

Voice Rebuild

Regenerate the selected span in the speaker's own voice — fix a damaged word or retake it with new text.

Voice Rebuild

Presets + ?

● Fast voice engine installed

✓ Pick a voice > **2 Type the text** > 3 Generate > 4 Place a take

VOICE REFERENCE Save Speaker...

Reference

From Selection

TEXT TO SPEAK From Transcript

What the take should say...

☐ Studio quality (slower)
Higher-quality multilingual cloning, ~2x slower than the fast engine

☒ Enhance Reference
Off keeps the original room/noise character in the new take

Language

English

Generate

The selected span (or a saved speaker) is the voice reference. After generating, a take replaces the selection, inserts at the cursor, or appends to the end. The take is fully synthesized speech — music or ambience under it is not preserved.

What it does. Voice Rebuild clones a voice from a short reference and speaks new text in it, entirely on-device. The selected span (or a saved speaker) is the voice reference; you type what the take should say, generate, then place the result. The output is fully synthesised speech, so any music or ambience under the original is not preserved — this rebuilds the voice, not the background. It runs on a fast English engine by default, with an optional higher-quality multilingual engine.

When to use it. Fix a clipped, coughed-over or mispronounced word in dialogue, replace a tongue-tripped phrase, or generate a short new line in the same voice without recalling the talent.

Controls

- **Voice Reference** — Choose *From Selection* (clones the selected span) or a previously saved speaker (no selection needed). **Save Speaker...** stores the current selection as a reusable voice; the trash icon removes a saved one.
- **Text to Speak** — What the new take should say. **From Transcript** prefills it with the transcribed words inside your selection (run Transcribe first).
- **Studio quality (slower)** — When the multilingual engine is installed, switches to higher-quality, multilingual cloning at roughly twice the time (default off).
- **Enhance Reference** — Cleans the reference voice; off keeps the original room and noise character in the take (default on).
- **Language** — Picks the language of the text (shown only with the multilingual engine installed; otherwise English).
- **Takes** — After generating, each candidate take is listed with a Play button to audition it, a word-accuracy and voice-match score, and the auto-chosen "best match" starred. The placement button (Replace / Insert / Append) shows where the take will land.

How to use it


1. Select clean speech as the voice reference, or pick a saved speaker.
2. Type the text, or pull it in with **From Transcript**.
3. Choose Studio quality and Language if available, and set **Enhance Reference**.
4. Click **Generate** and wait (you can Cancel a running job).
5. Audition the **Takes**, then click **Replace**, **Insert** or **Append** on the one you want — placement follows your current selection or cursor.

Tips

- The reference must be clean, isolated speech; pick a span with no music or overlap for the best match.
- Placement is decided when you click: a selection is replaced, a placed cursor inserts, otherwise the take appends to the end.
- Save a strong reference as a speaker so you can rebuild more lines later without re-selecting audio.
- Voice engines must be bundled with the app or installed locally before use; add the multilingual engine when you need more languages.

Prompt Isolate

Type what you want, and the AI pulls it out of the mix.



What it does. Prompt Isolate listens to your audio and separates out a single sound you describe in plain words — "acoustic guitar", "dog barking", "lead vocal". Under the hood it splits the audio into the described sound and everything else, then commits the side you asked for: **Isolate** keeps only the described sound, **Remove** subtracts it and keeps the rest. Everything runs locally on your computer. The separation engine is a one-time multi-gigabyte download; until it is present, the panel offers the download (with the model's license to acknowledge).

When to use it. Reach for it when a normal stem splitter has no category for what you need — isolating a specific instrument, lifting a sound effect, or surgically removing an unwanted noise (a phone ringing, a bark) while leaving the rest of the recording intact.

Controls

- **Described sound** — a text field where you type the sound to separate (for example, "acoustic guitar"). Be specific; this single phrase drives the whole separation. Required before you can run.
- **Quality** — how many passes the AI makes, then keeps the most consistent result. **Fast (1 pass)** is quickest, **Standard (3)** balances speed and reliability, **Best (6)** runs the most passes for the steadiest separation (default Standard).
- **Isolate** — replaces the audio (or selection) with only the described sound. If stems aren't ready yet, it separates first and then applies automatically in one step.

- **Remove** — subtracts the described sound from the audio (or selection), keeping everything else. Like Isolate, it separates first if needed.
- **Cancel** — replaces the buttons while a separation is running, so you can stop it.

How to use it


1. Select a region (or leave nothing selected to process the whole file).
2. Type the sound in **Described sound** and pick a **Quality**.
3. Click **Isolate** to keep only that sound, or **Remove** to take it out. The separation runs if needed, then the result is committed as one undoable edit.
4. Change your mind? Undo, and click the other button — the stems for the same prompt are cached, so the second choice applies instantly.

Tips

- Isolate and Remove are both undoable — try one, undo, try the other from the same cached split.
- Editing the audio or changing the prompt invalidates the cached stems; the next click separates again.
- Start on **Standard**; only reach for **Best** when a separation is leaking or unstable.
- Short, plain descriptions usually beat long ones — name the sound, not the scene.

AI Retouch

Describe how it should sound — and let the on-device AI restore and master the whole take.



What it does. AI Retouch is a generative, all-in-one restore-and-master pass that runs entirely on your computer. You type a plain-language instruction (or pick a preset chip) and the AI rebuilds the take to match — pulling down room reverb, repairing clipping and distortion, re-balancing the EQ, taming harshness, opening up squashed dynamics, and widening the stereo image, all in a single pass. Leave the instruction blank and it runs in Auto mode, applying a general automatic restoration. There is no separate preview or "Listen" residual here: Process applies the result directly to your audio as one undoable edit.

When to use it. Reach for it when a recording has several problems at once — a washy, distorted phone clip, an over-compressed dialogue line, a dull mix that needs life — and you would rather describe the goal than chain four modules. It is strongest on damaged material; on an already-clean track Auto can dull it, so use a Shape preset instead.

Controls

- **Instruction** — a free-text field where you describe how the audio should sound (for example "add warmth" or "reduce reverb"). Leaving it empty is valid and means Auto. The text field always overrides whichever preset chip is highlighted.
- **Presets** — a wrapping row of one-tap chips that fill the Instruction field for you, grouped as Restore (Auto restore, De-reverb, De-clip / distortion, Reduce muddiness, Fix mic / recording, Restore dynamics, Balance EQ) and Shape (Add clarity, Brighter, Add air, Darker, Add warmth, More bass, More punch, Vocals forward, Make louder, Widen stereo). Hover any chip for a short description; the active chip is outlined only while the text matches it exactly.
- **Prompt adherence** — how strongly the AI follows your instruction. Higher is more aggressive and can over-process (range 1–15, default 1).
- **Quality** — the number of refinement passes. Higher gives a cleaner result but takes longer (range 5–20, default 10).
- **Process / Cancel** — runs the engine; Cancel appears while it is working. A status row and "Local AI" badge confirm everything stays on-device. A "Credits & License" link opens the open-model credits for the engine.

How to use it


1. Select the region to treat, or select nothing to process the whole file.
2. Type an instruction, or tap a preset chip (or leave it blank for Auto).
3. Set **Prompt adherence** and **Quality** to taste.
4. Click **Process**. The result is applied directly and can be undone.
5. Not happy? Undo, tweak the wording or sliders, and Process again.

Tips

- Keep instructions to one clear change. Stacking several goals in one sentence can muddy the result and even push the AI the wrong way.
- The engine must be bundled with the app or installed locally before use. If the panel says the engine is not installed, run the installer or place the model files in the expected support folder, then use **Check Again**.
- Higher Prompt adherence plus higher Quality is the most assertive — but also the slowest and most likely to over-process. Start low and climb.
- On clean source, prefer a specific Shape preset over Auto, which can darken or dull an already-good track.

Generate

Make new audio from a description — drop in a clip, repaint a selection, or grow the file.



What it does. Generate creates brand-new audio on-device from a text prompt, with a choice of three engines. It works in shapes that follow what you have selected: with nothing selected it makes a fresh clip and inserts it at the playhead; with a time selection it repaints (inpaints) that region, conditioned on the audio around it so the new material blends in and the length stays the same; **Extend** grows a continuation out of the existing audio and appends it to the end; and with ACE-Step, **Cover** re-renders the whole file in a new style while keeping its melody. Everything runs locally, with no upload.

When to use it. Drop in room tone, a riser, rain, or a music bed where you have a gap. Repaint an unrepairable patch by replacing the selection with freshly generated audio that matches its surroundings. Extend a track that ends too soon — or Cover an entire song into a different style.

Controls

- **Engine** — *Stable Audio 3* (sound effects, ambience & short instrumental clips, 44.1 kHz), *ACE-Step XL Turbo* (full songs with vocals & lyrics, up to minutes long, 48 kHz — the default), or *Dasheng AudioGen* (whole audio **scenes** — intelligible speech, music, sound effects and ambience generated together in one coherent 16 kHz clip of up to 10 s). Each engine's weights download on demand.
- **Placement** — the panel reads your selection and playhead to choose the right action automatically. With no selection it generates at the playhead; with a time selection it inpaints that span; **Extend** grows a continuation from the end of the file.
- **Add to Selection** (*with a time selection*) — the alternative to repainting: generates a fresh clip from the prompt at the selection's length and mixes it **on top of** the existing audio. The original is kept underneath unchanged, and the added layer fades in and out at the edges so the joins never click. Needs a prompt (the layer is pure text-to-audio) and works with every engine that can generate the selection's length in one run — use it to lay rain over dialogue, slide a riser under an existing bed, or thicken a section without erasing it.
- **Prompt** — describe the sound (instruments, mood, effects, even tempo). Required for a fresh generated clip; optional for Inpaint and Extend, where leaving it blank fills from the surrounding audio. Press Return to run.
- **Lyrics / Speech** (*ACE-Step & Dasheng AudioGen*) — on ACE-Step, optional words to sing (one line per lyric line; blank for an instrumental). On Dasheng AudioGen the field becomes **Speech**: the exact words spoken in the scene, in any language — describe the voice itself (age, tone, style) in the Prompt. (Stable Audio 3 has neither, so the field is hidden there.)
- **Presets** — twelve one-tap prompt chips (Ambient pad, Rain & thunder, Vinyl crackle, Cinematic riser, Drum loop, Morning birdsong, Sub bass, Footsteps, Ocean waves, Acoustic guitar, Whoosh, Cafe ambience). Tapping one fills the Prompt field; your own text always wins.
- **Duration (s)** — length to generate (Stable Audio 3: 1–45 s; ACE-Step: 10–240 s; default 10). Longer takes proportionally more time. Locked to the selection length in Inpaint mode, and clamped into range when you switch engines. Dasheng AudioGen has no Duration slider — the model picks the natural length for the described scene (up to 10 s); timing words in the prompt ("brief", "a single knock", "sustained") influence it.
- **Context / Strength** — controls how much surrounding audio guides inpainting or generation at the cursor (0–15 s each side), and how strongly the regenerated material replaces the selected span (Strength appears for Inpaint only).
- **Cover strength** (*ACE-Step only*) — how far a Cover moves from the original (0–100 %, default 80 %). Lower stays close to the source; higher re-styles further toward the prompt/lyrics. Melody and timbre are preserved either way.
- **Quality** — a percent control that maps to the engine's generation-step window (4–12 steps on Stable Audio 3 / ACE-Step; 10–40 on Dasheng AudioGen, whose model-card default of 25 steps sits at 50 %). Higher values are cleaner but proportionally slower.

- **Guidance** — how strictly the result follows your prompt (0.5–4.0 on Stable Audio 3 / ACE-Step, default 1.0; 1–9 on Dasheng AudioGen, recommended 5.0). Higher hugs the prompt more tightly but roughly doubles generation time.
- **Rebuild highs (AI)** (*Dasheng AudioGen, default on*) — Dasheng renders 16 kHz-native scenes, so everything above its ~8 kHz content edge starts empty. This pass reconstructs the missing top band with the same spectral-recovery system the Spectral Recovery module uses (the AERO music net when installed, the model-free harmonic rebuild otherwise). Turn it off to keep the raw band-limited render.
- **Credits & License** — opens the model credits and license terms for the active engine (Stability AI for Stable Audio 3, ACE Studio for ACE-Step, Xiaomi MiLM Plus & SJTU X-LANCE for Dasheng AudioGen — the Dasheng sheet also carries the full prompt-tag reference).

How to use it

1. Pick the **Engine**: ACE-Step for songs, vocals and covers; Stable Audio 3 for effects and ambiences; Dasheng AudioGen for complete scenes that mix narration, effects, music and ambience.
2. Position the playhead for a fresh clip, make a time selection for Inpaint, or plan to use **Extend** / **Cover**.
3. Type a Prompt (plus Lyrics on ACE-Step) or tap a preset, then set Duration, Quality, and Guidance. Every run draws a fresh random take.
4. Click the action button (Generate / Inpaint Selection / Add to Selection / Extend / Cover). A progress bar shows the percentage; Cancel stops it.
5. The result is placed and selected automatically as a single undoable edit — there is no separate apply step. Press Undo if you do not like it, then simply re-run — on Auto every run rolls a fresh take.


Tips

- Start fast and cheap: Guidance 1.0 and lower Quality to audition ideas, then raise both once the prompt is right.
- Same prompt, different result: every run is a new take — when one lands, use Undo/Redo to compare takes, and save the one you keep.
- For seamless repairs use Inpaint and keep the prompt blank (or close to the surrounding sound) so the fill grows from context rather than introducing something new. To keep what is there and layer something over it instead, use **Add to Selection**.
- Cover works on the whole open file — describe the target style in the Prompt (and Lyrics, if sung) and use Cover strength to decide how far from the original to go.
- **Dasheng AudioGen structured captions** — a plain prompt is used as the scene description, but you can decompose the scene into the model's views by typing its tags directly in the Prompt: `<|caption|>` (overall scene — required first), then any of `<|speech|>` (voice style), `<|asr|>` (spoken words), `<|sfx|>`, `<|music|>`, `<|env|>`. Only include tags that are relevant. Example: `<|caption|> A gritty detective narrating over heavy rain. <|speech|> gritty deep male voice <|asr|> The city never sleeps. <|sfx|> heavy rain hitting pavement <|music|> melancholic solo saxophone <|env|> distant urban ambience`.

- Dasheng AudioGen writes all descriptive tags in **English** — only the spoken words (the Speech field / `<|asr|>`) use the target language. Keep Quality at 50 % (25 steps) and Guidance at 5.0 — the model-card defaults — and let the model choose the clip length.

Text to Speech

Type it, pick a voice, and place synthesized speech straight into the timeline.



What it does. Text to Speech synthesizes spoken audio from text you type, entirely on-device, and places the result into your file — replacing the selection, inserting at the cursor, or appending at the end. Four voice sources are available: fast preset voices, a natural-sounding engine with stronger pronunciation, a studio-quality multilingual engine, and your own cloned voices saved from Voice Rebuild. Each take is auditioned before you commit it, and placing it is a single undoable edit.

When to use it. Add a narration line you never recorded, patch a missing word or sentence in a voiceover, rough in scratch dialogue before the real session, or generate spoken content in another language or in a voice you've cloned.

Controls

- **Status row** — shows whether the selected engine's voices are installed, with an on-device tag ("MLX GPU" on macOS, "On-device" on Windows). If an engine is missing, an install hint appears with the location its model is expected in.

- **Voice source** — *Preset* (the on-device Kokoro model, fast English voices), *Chatterbox* (a natural engine with stronger pronunciation), *Qwen* (studio-quality and multilingual), or *Cloned* (speaks your saved voices through Voice Rebuild).
- **Accent / Language** — for *Preset*, an **Accent** picker (English (US) / English (UK)) that filters the voice list — Kokoro speaks English only. For *Chatterbox* and *Qwen*, a **Language** picker with ten languages (English, Chinese, French, German, Italian, Japanese, Korean, Portuguese, Russian, Spanish). Cloned voices speak English, so no picker is shown.
- **Voice** — the voice to speak with. *Preset* lists Kokoro voices for the chosen accent (marked ♀/♂); *Chatterbox* and *Qwen* clone a bundled preset reference clip; *Cloned* lists the speakers you saved in Voice Rebuild with **Save Speaker**.
- **Text to Speak** — the text field for what the voice should say.
- **Speed** — speaking rate from 0.50× to 2.00× (default 1.00×, the natural pace).
- **Generate / Cancel** — synthesizes the text in the chosen voice; Cancel stops a running synthesis.
- **Take row** — appears after a successful run: a **Play** button auditions the take, and the placement button — **Replace** (a selection exists), **Insert** (at the cursor), or **Append** (at the end) — commits it as one undoable edit.
- **Credits** — opens the model attribution and license details.

How to use it

1. Decide where the speech goes: make a selection to replace, park the cursor to insert, or neither to append at the end.
2. Pick a **Voice source** and a **Voice** (plus Accent or Language where offered).
3. Type the **Text to Speak** and set the **Speed**.
4. Click **Generate**, then use **Play** in the take row to audition.
5. Not right? Edit the text or voice and Generate again. Happy? Click **Replace** / **Insert** / **Append** to place it.

Tips


- Speech is fully synthesized — any music or ambience under a replaced selection is not preserved. Patch clean-voice regions, not mixed ones.
- *Preset* is the fastest for quick scratch lines; *Qwen* gives the most polished, multilingual results; *Cloned* keeps a consistent identity across your project.
- To speak in your own (or a speaker's) voice, first save it in **Voice Rebuild** with **Save Speaker** — it then appears under *Cloned*.
- Placement is undoable, so it's safe to place a take, listen in context, and undo if it doesn't sit right.

PART

Tools & Workflows

Batch Processing

When you have a folder full of files that all need the same treatment — a podcast season that needs de-humming, a stack of dialogue takes that need the same noise reduction, or a library you want to convert to a single format — Batch Processing applies your chosen modules to every file in one pass, unattended.



Open it from the **File** menu with **Batch Process...**, or press **⇧⌘B** (Ctrl+Shift+B on Windows). The batch window floats above the editor, so you can keep working while it runs, and you can drag it anywhere by its title bar. Closing it (the **×** in the corner) hides the window but keeps your list and settings; reopen it and everything is still there. You can't close it while a run is in progress.

Adding files

The left column is **Input Files**. Add files two ways:

- **Drag audio from Finder (macOS) or File Explorer (Windows)** straight onto the list. The drop area highlights when files are over it.
- Click **+** to open a file picker (you can select many files at once).

Only readable audio files are accepted, and duplicates are ignored automatically, so you can drop the same folder twice without piling up copies. The header shows a running count. **Trash** clears the whole list.

Building the pipeline

The middle column is the **Pipeline** — the ordered list of processing that runs on every file. To add a step, pick a module from the menu at the top and click **+**. Each step you add appears as a row with controls on it:

- The **power button** bypasses or re-enables that step without deleting it.
- The **sliders** button opens that step's full settings panel in the right-hand **Module Settings** column. Whatever you dial in there is the exact treatment applied to *every* file — there's no per-file tweaking.
- The **up/down chevrons** reorder the step (processing runs top to bottom).
- The **trash** button removes it.

You can stack as many steps as you like and they run in sequence — for example, De-hum, then a noise reduction, then a gain stage.

Running your saved Module Chain too

Below the pipeline is the **Module Chain** toggle. If you've built up a chain of modules in the main editor, this option also runs that whole chain on each file, *after* the batch pipeline steps above. The toggle shows how many steps it holds and lists them; it's disabled until you've actually added modules to the editor's chain. Use this when you've already perfected a chain on one file and want to apply the identical sequence across the batch.

Output: where, what format, and naming

The **Output Folder...** button sets the destination; the folder name appears next to it. A destination is required before you can run.

Next to it, choose the **output format** for every file — the same format choices as the app's Export panel (WAV, AIFF and CAF at various bit depths, plus the on-device-encoded compressed formats). A **sample rate** menu sits beside the format ("Original" keeps each file's own rate; picking a rate converts every file to it), compressed formats add a **Bitrate** (or **Quality**) menu, and PCM formats offer a **Dither (16/24-bit)** option. Source files are converted to this one format on the way out.

Rename Output wraps every saved file in an optional **Prefix** and **Suffix** — handy for tagging results like `clean_take01_processed`. A live example filename updates as you type so you can see exactly what will be written.

Existing files are never silently overwritten. If your chosen names would collide with files already in the destination, a warning lists them before the run starts; choose **Continue** and the new outputs are saved with a numbered suffix (such as `name-2`) so nothing is lost, or **Cancel** to rethink.

Running and monitoring

Click **Run Batch** (it stays disabled until you have at least one file and an output folder). A progress bar and a `Processed X / Y` readout appear, the footer reports how many files are queued, and each file in the list gets a green check or red mark as it finishes. A file that fails shows the reason inline beneath its

row — and again as a tooltip — so a failure tells you *why* rather than leaving you with a bare red mark; the completion line likewise counts the failures and points you at those per-row reasons. Hovering a file that succeeded shows the name it was written under.

The run is also mirrored on the app's shared progress panel, so you can keep an eye on it with the Batch window pushed aside. **Cancel** — the Batch window's own button, or the shared panel's — stops cleanly after the file currently being processed; finished files are kept, the rest are skipped, and the readout tells you how many completed.

Once a run has finished, a **Show in Finder** button (**Show in File Explorer** on Windows) appears beside **Run Batch**. It reveals the files that run actually wrote — or, if nothing was written, the output folder itself — and points at the folder the run used even if you have since chosen a different one.

Which modules are available — and why some aren't

The module menu only lists modules that can run start-to-finish on a whole file without you in the loop. Three families are intentionally left out: modules that need a hand-drawn spectral selection, gap, or hosted plug-in (such as Spectral Repair, Interpolate and Plugin Host); analysis and generation tools that produce a sidecar result or need a live prompt rather than transforming the file (such as Transcribe, Stem Remix, Prompt Isolate and Generate); and modules that change a file's length or channel count (such as Time & Pitch, Resample, Channels and Stereo Width). They simply can't be replayed safely across many files unattended.

One reference-based module *is* offered: **Match**. It runs over a batch only in its EQ-curve mode and only with a reference already learned — each file is re-analyzed against that stored reference. A Match step without a learned reference (or one saved in Full Master mode) can't run; it is flagged orange in the pipeline list ("will be skipped"), and the footer counts how many steps won't run, so you know before you press **Run Batch**. Everything else offered in the menu is guaranteed to do real work on each file.

Markers

Markers are named flags pinned to exact points on the timeline. Use them to label takes, mark edit points, jump between sections, or keep a list of moments you want to return to. They sit on the time ruler above the waveform and spectrogram as small green flags, each with its name shown next to it.



Markers are saved as part of your editing session and move with your audio: when you cut, paste, trim or resample, the markers shift so they stay attached to the same moment in the sound. Adding, naming, deleting and clearing markers are all undoable.

Adding a marker

A marker is always placed at the playhead — the vertical line showing the current position.

- **Add Marker** (`Markers & Regions` ► `Add Marker` , shortcut `⌘⇧M`, or `Ctrl+Shift+M` on Windows) drops a marker at the playhead. New markers are auto-named "Marker 1", "Marker 2", and so on, in the order you create them.

To place a marker precisely, click or scrub to set the playhead first, then add the marker. You can also add one from the Markers List window with **Add at Playhead** (see below).

Navigating markers

Markers give you fast jumping between points without hunting on the timeline.

- **Next Marker** (`Markers & Regions` ► `Next Marker` , shortcut `⌘→`, or `Ctrl+Alt+→` on Windows) moves the playhead to the first marker after the current position.
- **Previous Marker** (`Markers & Regions` ► `Previous Marker` , shortcut `⌘←`, or `Ctrl+Alt+←` on Windows) moves the playhead to the nearest marker before the current position.

These commands are greyed out when there is no marker in that direction. In the Markers List, the small play icon beside each marker also jumps the playhead straight to it.

The Markers List

Open **Markers & Regions ▶ Markers List...** for a window that shows every marker in the file, with its name and its position as minutes-and-seconds. From here you can:

- **Rename** a marker by clicking its name and typing. The new name updates everywhere, including the flag on the ruler.
- **Jump** to a marker with the play icon at the right of its row.
- **Delete** a single marker with the trash icon at the right of its row.
- **Add at Playhead** to drop a new marker without leaving the window.
- **Export CSV...** to save the list (see below).

The window title shows the current marker count. If there are none yet, it reminds you that ⌘M (Ctrl+Shift+M on Windows) drops one at the playhead.

Deleting and clearing

- **Delete Marker** (**Markers & Regions ▶ Delete Marker**) opens a submenu listing every marker by name; pick one to remove it. This is handy when you want to delete by name rather than by position.
- **Clear All Markers** (**Markers & Regions ▶ Clear All Markers**) removes every marker at once.
- In the Markers List, the trash icon on any row deletes that one marker.

All of these are undoable, so a mistaken clear can be reversed.

Markers from Transcribe

The **Transcribe** module can turn speech into markers automatically. In its panel, turn on **Add markers for segments** before transcribing, and you get one marker at the start of every spoken segment, named with the words of that segment (truncated if very long). If you have already transcribed, the **Create Markers** button does the same thing for the existing transcript; it changes to **Markers Added** once the markers exist. Re-running it is safe — it won't pile up duplicates at the same spot.

A related tool, **Find Similar Events** (**Markers & Regions ▶ Find Similar Events** , shortcut ⌘E, or Ctrl+Alt+E on Windows), takes your current selection as a template, scans the whole file, and drops a marker at each similar moment it finds, each named "Similar" with its match percentage. This needs a time selection that is shorter than half the file, so there's material left to search.

Exporting markers

From the Markers List, **Export CSV...** writes a comma-separated file containing each marker's name, its position in seconds, and its exact frame position. The file is named after your document by default. This is a plain spreadsheet-friendly file you can open in any spreadsheet app or feed into other tools. (Transcribe also offers a separate **Export SRT...** for subtitle files, covered in the Transcribe chapter.)

Regions

Where a marker pins a single *point*, a **region** is a named time *range* — a labelled span you can recall, play, and export as its own file. Regions draw directly on the waveform and spectrogram as violet spans with a solid border and a small name tab — one fixed region tint, deliberately distinct from the yellow dashed live selection so a region can never be mistaken for it. Regions share the **Markers & Regions** menu, and like markers they follow the audio through cuts, pastes, trims and resamples — with one extra rule: a region whose span collapses to nothing in a cut is dropped. Creating, renaming, deleting and clearing regions are all undoable.

- **Create Region from Selection** (**Markers & Regions** ▶ **Create Region from Selection** , shortcut ⇧⌘R, or Ctrl+Shift+R on Windows) turns the active time selection into a region. New regions are auto-named "Region 1", "Region 2", and so on. Right-click a region on the waveform or spectrogram to **Rename** or **Delete** it in place.
- **Next Region / Previous Region** jump the playhead between region starts, mirroring the marker navigation (these two have no default shortcut).
- **Recall**. Selecting a region re-applies its span as the active time selection and parks the playhead at its start — the headline difference from a marker, which only seeks.

The Regions List

Markers & Regions ▶ **Regions List...** opens a window showing every region with its name (click to edit — the rename is undoable), start and end times, and duration. Each row has buttons to **play** just that region, **select** its range in the editor, or **delete** it. The window also offers **Create from Selection**, **Export CSV...**, **Import CSV...**, and **Export as Files...** without leaving it.

Deleting and clearing

- **Delete Region** (**Markers & Regions** ▶ **Delete Region**) opens a submenu listing every region by name; pick one to remove it.
- **Clear All Regions** removes every region at once. Both are undoable.

Exporting regions as files

Markers & Regions ▶ **Export Regions as Files...** renders each region to its own audio file — the fastest way to turn one long take into named deliverables. A panel lists the regions with per-item include checkboxes, a **naming pattern** built from `{name}` (the region's name), `{index}` (its number) and `{file}` (the document name) with a live filename preview, a destination folder, and the same format options as the Export dialog (format, sample rate, bitrate/quality, dither). The export runs with a progress bar and a working Cancel, and reports any per-file failures.

Markers & Regions ▶ **Split by Markers...** is the same panel driven by your *markers* instead: the file is split into the segments between consecutive markers, including the leading audio before the first marker and the tail after the last one. Each segment is named after the marker that opens it (the leading segment has no marker, so it falls back to a numbered "Segment" name). Uncheck any segments you don't want before exporting.

CSV import

Both lists round-trip through CSV. **Markers & Regions ▶ Import Markers CSV...** reads `name, seconds, frame` rows — exactly what the Markers List's Export CSV writes — and **Markers & Regions ▶ Import Regions CSV...** reads `name, start_seconds, end_seconds, start_frame, end_frame`; a header row is optional in both. Each import lands as a single undoable step and reports how many rows were imported and how many were skipped.

Markers, regions, and the saved file

When you save to **WAV or AIFF**, markers and regions are written into the file itself as standard embedded metadata: markers become **cue points** (WAV cue/label chunks, AIFF MARK entries) and regions become labelled ranges. They survive a full round trip — save the file, reopen it (in FFT or in any editor that reads cue points), and the markers and regions are still there. Opening a file that already carries cue points imports them as markers and regions in the same way. Any other embedded metadata the file carries (Broadcast Wave `bext`, iXML, ID3 and similar) is preserved untouched across a save, as long as the destination stays in the same container family (WAV to WAV, AIFF to AIFF).

Statistics & Analysis

The **Statistics** panel gives you an instant, measured read-out of your audio: how loud it is, how close to the ceiling it sits, whether it clips, and whether it carries an offset. It answers the questions you can't reliably judge by eye or ear alone — "Am I clipping?", "Is this at broadcast loudness?", "Is there DC on the file?" — with hard numbers.

Statistics	Whole file
Duration	0:03.000
Channels	2
Sample peak	-6.0 dBFS
True peak	-6.0 dBTP
RMS	-9.8 dBFS
Integrated	-7.6 LUFS
DC offset (max ch)	-0.00000
Clipped	0 samples

Opening the panel

Click the **bar-chart icon** (its tooltip reads "Waveform statistics") in the editor toolbar. A small popover opens beneath it. The button is dimmed until a file is loaded.

Analysis runs the moment the popover opens, in the background, so the app stays responsive even on long files. While it works you'll see **"Analyzing..."** with a spinner; the numbers appear as soon as the scan finishes. If you edit the audio while the panel is open, it quietly re-measures so you never see figures from a stale version of the file.

Selection vs. whole file

The header tells you exactly what's being measured:

- **"Selection"** — you have a region selected (a time range on the waveform, or a spectrogram selection), so every figure describes *only* that region.
- **"Whole file"** — nothing is selected, so the figures describe the entire file.

This makes the panel a quick way to compare two parts of a recording, or to check a single phrase, a noise-only gap, or a problem spot against the rest of the track. To switch between the two, just change your selection (or clear it) and reopen the panel.

What each metric means

- **Duration** — the length of what's being measured, shown as minutes:seconds.milliseconds (hours appear too if the range is long enough).
- **Channels** — how many channels are in the file (1 for mono, 2 for stereo, and so on).
- **Sample peak** — the loudest single sample, in dBFS. **0.0 dBFS** is the digital ceiling; anything above it can't be stored and means clipping. Healthy masters typically peak a little below 0 (for example around **-1 dBFS**) to leave headroom.
- **True peak** — the *real* peak of the waveform between samples, in **dBTP** (decibels true peak — the BS.1770 unit, with the same 0 = full-scale reference as dBFS). Because the continuous signal can crest higher than any stored sample, true peak is often a touch louder than sample peak. This is the figure that matters for delivery: streaming and broadcast specs usually ask for a true-peak ceiling around **-1 dBTP**. If true peak reads above 0, your file may distort after format conversion even when the sample peak looks safe.
- **RMS** — the average level, in dBFS, which tracks perceived loudness better than a single peak. There's no single "correct" value; it's most useful for comparing the density of one passage against another.
- **Integrated** — integrated loudness in **LUFS**, the standard loudness measurement used for streaming and broadcast. It's gated so silences don't drag the figure down. Common targets are around **-14 LUFS** for streaming and **-23 LUFS** for broadcast — check your delivery platform's spec.
- **DC offset (max ch)** — any constant bias in the waveform, where it should be centred on zero. The reading is the single worst channel's bias (the largest by magnitude, with its sign), so a one-channel converter fault can't hide in a cross-channel average. A reading near 0 is healthy; a value noticeably away from zero wastes headroom and can cause clicks at edits. If you see one, the **Channels** module offers a way to remove it.
- **Clipped** — the count of samples sitting at or beyond the digital ceiling. **Zero is what you want.** Any non-zero count is shown in **red** as a warning that the audio is already distorted; consider the **De-clip** module to repair it.

When a value is silent or empty (for example measuring a gap), level figures read **-∞ dBFS** or **-∞ LUFS** rather than a number — that's expected, not an error.

Tips

- For loudness work, trust **Integrated (LUFS)** for overall loudness and **True peak** for your ceiling — not Sample peak alone.
- A red **Clipped** count means damage already exists in the file; fix it before mastering, not after.
- Use the **Selection** mode to spot-check a quiet passage's noise floor (its RMS) versus the loud parts.
- True peak and LUFS over a very long whole-file scan take a moment longer to compute — that's the background analysis being thorough.

The Notes Editor (Pitch & MIDI)

The Notes view turns your audio into editable notes on a piano roll, so you can retune, retime, split, level, and reshape a performance the way you would in a Melodyne-style pitch editor. Open it with **View ▶ Notes** (⌘4, or Ctrl+4 on Windows). Detection runs automatically when the view opens (you'll see **Detecting notes...** while it works), and again whenever the audio changes. Everything you do previews live and stays non-destructive until you leave the view and choose to keep it.



The right end of the control strip shows how many notes were found, or how many are selected.

Detection

Detection is **monophonic** — it hears one voice at a time, and is at its best on vocals, a solo line, or a single instrument. Chords and dense mixes may detect nothing usable; if that happens you'll see a "No notes detected" hint. The detector is bundled with the app — nothing to download.

If a detection has gone off the rails, **View ▶ Re-detect Notes** re-runs it from scratch — it asks first, because it discards every detection correction and pending edit.

Edit and Assign

The **Notes** switch chooses what your gestures act on:

- **Edit** corrects the *tune*: pitch, timing, length, level, vibrato — the audible edits. While you work, playback auditions your edits in context.
- **Assign** (Note Assignment) corrects the *detection*: splits, glues, octave and pitch relabelling — the note map that Edit then works on. Assign changes are metadata only and never touch the sound; playback in Assign always auditions the **original** audio. That's why deleting a note in Assign sounds like nothing happened — you've told FFT the note was mis-detected, not asked it to remove the sound. To silence a note audibly, delete it in **Edit**.

Assign edits keep their own local Undo/Redo (⌘Z / ⇧⌘Z, or Ctrl+Z / Ctrl+Shift+Z on Windows inside the Notes view, plus the **Undo/Redo** buttons), separate from the audio undo history.

The tool palette

The **Tool** picker (icons in the control strip) chooses what a click or drag does; the cursor changes to teach you each tool. Each tool also has a single-letter keyboard shortcut — **A** Main, **S** Separate, **V** Vibrato, **D** Drift, **I** Intensity, **G** Glide — that switches tool the instant you press it while the piano roll has focus (a letter only fires when that tool is offered for the current mode, and types into a note-name field when one has focus). Which tools appear depends on whether you're in Edit or Assign:

- **Main** (arrow) — always available. Drag a note up or down to retune it; hold ⌘ (Alt on Windows) while dragging for free, un-snapped pitch. Drag a note's left or right **edge** to change its length. In Edit, double-click a note's body to snap it to pitch.
- **Separate** (scissors) — a single **click** on a note cuts it where you click (the cut-line preview shows where); double-click a separation line to glue two notes back together; drag a line to move the boundary. Cuts snap to the playhead when it's nearby, then to a note onset, then to a zero crossing, for clean, click-free edits.
- **Vibrato** (Edit) — drag up to intensify a note's pitch modulation, down to flatten it (and past flat to invert); double-click to restore the original.
- **Drift** (Edit) — the same gesture for slow pitch drift (scoops and sags); double-click to restore.
- **Intensity** (Edit) — drag a note up or down to change its level; double-click to reset it to 0 dB.
- **Glide** (Edit only) — shapes the portamento across the soft join between two notes. Drag the join itself left/right to shorten or lengthen the glide — no pre-selection needed — or select two or more adjacent notes and drag up/down to shape every soft join between them at once. Double-click a join to reset it. An amber connector draws the resulting pitch trajectory.

In Assign, only **Main** and **Separate** are offered — the others are audible edits, which is Edit's job. There is deliberately no separate "free pitch" tool: whether a drag lands on the pitch grid is the **Quantize** toggle, and holding ⌘ (Alt on Windows) while dragging with the Main tool momentarily inverts it.

Selecting, retuning, and retiming

Click a note to select it; Shift- or ⌘ -click (Ctrl-click on Windows) to add notes; drag an empty area to marquee-select; ⌘A (Ctrl+A on Windows) selects every detected note at once. With one or more notes selected, in Edit mode:

- **Drag vertically** to retune. With **Quantize** on, the move snaps to the chosen **Scale** (or to the nearest semitone if the scale is Chromatic); with it off, pitch is continuous. Hold ⌘ (Alt on Windows) while dragging to momentarily invert the Quantize toggle. A multi-note drag stays parallel, and the dragged note auditions as a soft tone so you can hear the target.
- **Drag horizontally** to shift notes in time.
- **Drag the edges** to lengthen or shorten.

Each drag is a single ⌘Z (Ctrl+Z on Windows) undo step, and **Escape** cancels a drag in flight.

Pitch correction

Set a **Scale** root and type (Chromatic, Major, Minor, pentatonics, Blues, modes). **Snap to Scale** moves the selected notes (or all of them) onto the nearest scale degree, with **Strength** controlling how strongly the correction is applied.

Separations: soft and hard

Adjacent notes show a thin line at their join. Besides the Separate tool's cut/glue/move gestures, **⌘-double-click** (Alt-double-click on Windows) a separation toggles it between **soft** (a smooth, glidable join) and **hard** (square brackets — fully independent notes). A hard join has no glide; some boundaries that were detected as hard can't be softened.

Correcting the detection (Assign)

In **Assign** the detection edits are the separation gestures: cut a note with the **Separate** (scissors) tool — a single click cuts where you click — double-click a separation line to glue, and drag a line to move the boundary. A small **Re-detect** button offers the same escape hatch as the View menu.

Turn on **Range** to work within a pitch band: drag the dashed blinds in the piano roll or type the low/high bounds. In Edit the band limits how far retunes can go; at detection, sound outside the band is folded into the in-band notes.

Zoom and the keyboard rail

A plain scroll or pinch zooms and scrolls the shared timeline, just like the waveform. Hold **Shift** to move the *pitch* axis instead: Shift-scroll pans the pitch view up and down, and Shift with **⌘** or **⌥** (Ctrl or Alt on Windows) (or Shift-pinch) zooms the pitch range. The piano-keyboard rail on the right edge shows where you are — drag it to pan, scroll or pinch on it to zoom, and double-click it to snap back to an auto-fit of the detected notes.

Applying your edits

There is no Apply button: edits are committed when you leave the Notes view. Switch to any other view with pending edits and FFT asks "**Keep your note edits?**" — **Keep** renders them into the audio (a single undoable step), **Discard** reverts to the original, **Cancel** stays in the editor. While you're still in the view, the pinned **Reset** button discards the live edits at any time.


Exporting per-note audio

- **Export Notes...** (the **Export** button in the Edit strip) saves each detected note as its own WAV file into a folder you choose, carrying your live Edit-mode corrections — handy for building sample sets from a performance.
- **File ▶ Export MIDI...** is tied to polyphonic note detection and is currently disabled while the editor detects one voice at a time.

Detection quality follows the material: a clean, dry solo line detects best; heavy reverb, chords, or dense mixes can defeat it.

The Text Editor

The Text view edits speech the way you edit a document: it transcribes the recording, you correct the words, and FFT re-speaks the changed sentences in the speaker's own voice. Open it with **View ▸ Text** (⌘5, or Ctrl+5 on Windows). Like the Notes editor, edits stay non-destructive while you work and commit when you leave the view.



Opening the Text view

When the view opens, FFT transcribes the audio with Whisper (you'll see **Transcribing...**) on **Multilingual — Auto-detect**. Text mode always starts there for a new file, whatever the **Transcribe** module's **Language** happens to be set to. If the transcription comes back empty, the **No speech found** card carries a language menu: pick the spoken language there and press **Re-transcribe**. That pick is remembered for this file only — open another file and Text mode is back on Auto-detect, and the Transcribe module's own Language is never changed by it. A Whisper transcription model must be installed; if none is, the view shows an install prompt instead (see the AI models chapter). The toolbar gains two Text-mode extras: a **Transcription** button that opens the Transcribe module (its own transcription run, with its own Model and Language settings), and a small whole-file waveform **locator** that marks where the currently selected transcript text sits in the audio.

Editing the transcript

The transcript is a list of sentences, each with a timestamp. Click a timestamp to **play** that sentence — its span is selected in the waveform and the locator, and playback runs from its start to its end (with Loop latched, it loops). Click into any sentence and type to correct it; edited sentences are highlighted and counted in the header ("2 edited"). Press **Return** (or Escape, or click anywhere else) when you're

done: the sentence is committed and the keyboard is yours again, so **Space** plays and pauses as everywhere else in FFT. Tab and Shift-Tab move editing to the next or previous sentence. Deleting all the text of a sentence marks that span for removal from the audio.

Each edited row offers a **restore** button (back to the original words), and the header's **Reset** discards every text edit at once. The header also carries a **Re-transcribe** button, which re-runs Whisper from scratch (it confirms first, since that discards your text edits and any rebuilt takes), and an **Export** menu that saves the current transcript to a file — **Plain Text (.txt)**, **SubRip (.srt)**, or **WebVTT (.vtt)** subtitles. The audio spans to regenerate are derived automatically from the transcript's word timestamps — you never make the selections yourself.

Rebuilding the edited sentences

Rebuilding is two steps, and the footer names the one that comes next:

1. **Rebuild N sentences**, at the foot of the transcript, re-speaks every edited sentence with **Voice Rebuild**, cloning the voice from the recording itself and speaking your corrected text. A single row's  button does the same for just that sentence. Each rebuilt row then gets a **play** button so you can listen to the new take in place, and its  becomes **try a different take**.
2. **Apply N rebuilt sentences to audio** writes the takes into the recording as one undoable edit. The transcript stays put — the applied rows now simply read what you typed, marked with a green dot — and the first applied sentence is selected, so Space plays the result right away. Nothing touches the audio before Apply.

The toolbar carries the synthesis-voice settings next to the search field: a **Rebuild voice** language menu, a **Studio** toggle for the higher-quality multilingual engine (slower, and it needs the studio voice model installed), and **Clean voice**, which denoises the voice reference before cloning — useful on noisy recordings.

Self-verification

Every generated take is checked: FFT re-transcribes the synthesized audio with Whisper and compares the words against what you typed. The dot at the end of each row shows the verdict — green or blue for a verified take, orange when Whisper heard different words, with the header summarising **Ready to apply** or **Check flagged segments**. A flagged take isn't discarded; listen to it, and if it really is wrong, try a different take.

Leaving the Text view

Rebuilt takes you haven't applied yet are applied silently when you switch away — the waveform simply shows the changes. If some edited sentences haven't been rebuilt yet, FFT asks "**Apply your remaining text edits?**" — **Apply** rebuilds and applies them, **Discard** drops the un-rebuilt edits (keeping the rebuilt ones), **Cancel** stays in the editor. Either way the result is a normal, undoable audio edit.

Searching the transcript

The search field above the transcript finds words as you type, shows the match count, and steps through matches with the arrows (or Return). Jumping to a match scrolls the transcript, selects the matched words in the waveform, and moves the playhead — a fast way to locate a phrase in a long recording.

Good to know

- Text editing works on speech; it won't transcribe or rebuild music.
- Everything runs on-device: Whisper for transcription and verification, Voice Rebuild for synthesis. Both models are managed from the AI models installer.
- A rebuilt sentence is a synthesis in the speaker's cloned voice — very close, but pacing and timbre can differ slightly from the original delivery. Audition before applying, and reroll takes you're not happy with.
- Editing the audio elsewhere invalidates pending takes: they're regenerated against the new audio the next time you rebuild.
- Applying keeps your transcript: the applied sentences are updated in place rather than the whole file being transcribed again. Undoing an Apply restores the audio; the transcript is then re-transcribed the next time the Text view looks at it.

PART

Exporting

Saving & Exporting

When you're ready to commit your work to disk, FFT gives you a single, focused export dialog that handles everything from a quick overwrite of the original file to a fresh render in a different container, bit depth, or codec. This chapter covers saving in place, the full Save As / Export dialog, exporting just a selection, and every option you'll find there.

Saving in place

- **Save** (, or Ctrl+S on Windows, or the down-into-tray button in the toolbar — the upward-arrow button beside it is **Export**) writes your edits straight back to the file you opened, in its original format. If FFT knows the file's format, it simply overwrites it. If it can't (for example, a brand-new untitled recording, or a format it can't write back to), Save automatically opens the full Save As / Export dialog instead.
- If you try to close a tab, open another file, or quit with unsaved changes, FFT asks whether to **Save**, **Don't Save**, or **Cancel** first, so you never lose work by accident.

The Save As / Export dialog

Open it with **File ▶ Save As / Export...** (, or Ctrl+Shift+S on Windows). This is FFT's own dialog, not the plain system save sheet, and it has three parts: a destination, a format, and format-specific options.

Choose where it goes. The **Destination** field shows the full path of the file to be written. You can type or paste a path directly, or click **Browse...** to pick a folder and name with the standard system file picker. The line beneath the field confirms which folder you're writing into. The extension is set for you to match the format you pick, so you don't have to type it.

Pick a format. The **Format** menu lists every container and encoding FFT can write:

- **WAV** — 8-, 16-, 24-, 32-bit PCM, 32-bit float, plus μ -law and A-law
- **AIFF** — 8-, 16-, 24-, 32-bit PCM
- **CAF** — 16- or 24-bit PCM, 32-bit float, and IMA4 ADPCM
- **FLAC** — 16- or 24-bit (lossless, compressed)
- **Apple Lossless (.m4a)** — 16- or 24-bit (on Windows the encoder is built into the app: mono/stereo, sample rates up to 65535 Hz — choose CAF or FLAC to keep a 88.2/96/176.4/192 kHz master lossless there)
- **AAC (.m4a)**, **MP3**, **Opus**, and **Ogg Vorbis** — lossy, compressed (on Windows, AAC is offered when the system's Media Foundation AAC encoder is present, which the N/KN editions lack)
- **AU/NeXT**, **Wave64**, and **RF64** — for specialist and very-large-file workflows

The MP3, Opus, and Ogg Vorbis encoders are built into the app and run entirely on your computer — nothing is uploaded and no extra software is needed. When one of these three is selected, a small **Encoded on-device · Credits & Licenses** link appears; click it to view the encoder credits.

The format already carries the bit depth. Because each entry names its own bit depth (for example *WAV (24-bit PCM)* versus *WAV (32-bit float)*), choosing the format is how you choose the depth — there's no separate bit-depth control.

Dither

When you export to **16- or 24-bit PCM** (in any container that supports it), a **Dither** option appears with the checkbox **Add TPDF dither**. Dither adds a tiny, inaudible amount of shaped noise just before the audio is reduced to integer samples, which smooths out the quantization steps and prevents low-level distortion on quiet fades and tails. Leave it on when you're producing a final 16-bit master from higher-resolution audio. It has no effect on 32-bit, float, or lossy formats, so the option simply hides itself for those.

Bitrate and quality

For the lossy formats, a **Bitrate** (or, for Ogg Vorbis, a **Quality**) menu appears so you can trade file size against fidelity:

- **AAC** — 96 to 320 kbps
- **MP3** — 128 to 320 kbps
- **Opus** — 64 to 256 kbps
- **Ogg Vorbis** — quality steps Q0 to Q10

Higher numbers mean better quality and larger files. These choices only appear for the format they apply to.

Sample rate

A **Sample rate** menu sits below the format. **Original** — the default — keeps the project's current rate untouched, with no conversion anywhere in the path. Pick one of the offered rates (44.1, 48, 88.2, 96, 176.4 or 192 kHz) to convert on the way out instead. The menu only offers rates the chosen format can actually encode — MP3 caps at 48 kHz and AAC at 96 kHz — and if you switch to a format that can't produce the rate you had picked, the choice falls back to Original rather than failing at export time.

The export menu is the quick way to *deliver* at a different rate; for a rate change you want to keep working on, or where you want control over the converter, use the **Resample** module before exporting.

Exporting a selection

To render only a highlighted region instead of the whole file, make a time selection on the waveform or spectrogram, then choose **File ▶ Export Selection...** (⌘⇧E, or Ctrl+Shift+E on Windows). The same dialog appears, with the same format and option controls, and only the selected audio is written to the new file. This item is available only while a selection is active. A large selection export runs in the background like any other big write (see *Finishing the export* below), and when it completes FFT reveals the new file in Finder (macOS) or File Explorer (Windows) so you can go straight to it.

Dragging a selection out

The fastest export needs no dialog at all. While a time selection is active, the **selection readout strip** (the Selection/View bar beneath the editor) shows a small **drag-out grip** next to the selection fields. Drag it to Finder (macOS) or File Explorer (Windows) — or to any app that accepts audio files — and the selection is written there as its own file, named after your document (for example `MyTake selection.wav`). The file is rendered only at the moment you drop it, so the drag itself is instant even on long selections.

Right-click the grip to choose the format it exports: WAV 16-bit, 24-bit or 32-bit float, or FLAC 16- or 24-bit. The choice persists between sessions; the default is 24-bit WAV. The grip is deliberately lossless-only — for lossy formats, sample-rate conversion, or dither, use the Export dialog instead.

Finishing the export

Click **Export** to write the file (the button is disabled until you've chosen a destination), or **Cancel** (`Esc`) to back out. Writes are atomic and safe: FFT renders to a temporary file and only swaps it into place once the write fully succeeds, so an interrupted export — a full disk, for instance — can never corrupt or destroy an existing file. If a file with the chosen name already exists, you'll be asked to confirm before it's replaced.

Large files are written in the background. A small file lands the instant you click **Export**, but for a big write — a long streamed file, or a multi-gigabyte project — FFT hands the encode off the main thread so the editor stays responsive, and a progress panel appears (*Saving ...* for a save, *Exporting ...* for a selection). Because the write is one continuous operation, the panel shows activity and elapsed time rather than a percentage. When the file is large enough to be *streamed*, that panel carries its own **Cancel** button — distinct from the dialog's **Cancel** above — and cancelling aborts the write cleanly, leaving no partial file behind. Only one save or export of a given document runs at a time: start a second while the first is still writing and FFT tells you to wait for it, or cancel it from the progress panel, rather than quietly dropping your request.

An emptied document can be saved, too. If an edit clears the whole file (Select All, then Delete), **Save** and the export dialog still write the — now zero-length — result out rather than refusing it.

Metadata travels with the save

When you save or export to **WAV or AIFF**, the file's embedded metadata is preserved: your markers are written back as standard cue points and your regions as labelled ranges (so both survive a round trip through FFT or any editor that reads cue points), and opaque metadata the source file carried — Broadcast Wave `bext`, iXML, ID3 and similar — is re-emitted untouched, as long as the destination stays in the same container family (WAV to WAV, AIFF to AIFF). See the Markers chapter for the details.

Other ways out of FFT

- **Markers & Regions ▶ Export Regions as Files...** and **Markers & Regions ▶ Split by Markers...** batch-render labelled sections of the file to their own files, with a naming pattern and the same format options as this dialog — covered in the Markers chapter.
- **File ▶ Export MIDI...** writes the Notes editor's detected notes as a standard MIDI file — covered in the Notes & MIDI chapter.

- **File ▶ Export Lyrics & Chords...** writes a classic chord-over-lyrics chart of the open file as a PDF, named after your document and ending in *Lyrics & Chords.pdf*. FFT isolates the lead vocal, transcribes it, reads the chord progression, key and tempo off the mix, and typesets the result; the job runs in the background like a large export. It needs the Voice Isolate and transcription models — if either is missing you are offered the downloads first. A **Lyrics language** choice (Automatic, or one of nine languages) pins the transcription when auto-detection struggles: on macOS it sits in the save dialog, on Windows it is asked just before the save dialog opens. On Windows the chart's bundled fonts do not yet include Chinese, Japanese or Korean glyphs, so lyrics in those languages come out as empty boxes in the PDF.
- **File ▶ Export Raw Data...** writes headerless samples for specialist workflows, mirroring **Import Raw Data...**

Tip: For archival masters, choose a lossless format — WAV, AIFF, CAF, FLAC, or Apple Lossless — and keep a 24-bit or float copy. Reserve MP3, AAC, Opus, and Ogg Vorbis for delivery, since each re-encode of a lossy file loses a little more quality.

PART

Reference

Keyboard Shortcuts

FFT maps the commands you reach for most to the keyboard, so you can keep your hands on the keys and your eyes on the waveform or spectrogram. This chapter is a complete reference, grouped by area. The shortcuts work on whichever editor window is in front; if a command is dimmed in its menu (for example, because there is no audio loaded or nothing selected), its shortcut is inactive too.

Every table lists the macOS chord and its Windows equivalent side by side. On macOS, **⌘** is Command, **⇧** is Shift, **⌥** is Option, and **⌃** is Control; on Windows the same commands use **Ctrl** where macOS uses **⌘** (or **⌃**), **Alt** where it uses **⌥**, and **Shift** for **⇧** — so **⇧⌘M** on macOS is **Ctrl+Shift+M** on Windows. A handful of Windows chords are deliberately different, because the direct translation would collide with another command or cannot be typed on some keyboard layouts; the tables show the chord the Windows menu bar actually advertises, and the *Windows notes* at the end of this chapter list the extras.

Transport and playback

Action	macOS	Windows
Play / Pause	Space	Space
Record from input	⌘R	Ctrl+R
A/B Compare Last Edit	⌥B	Alt+B

Stop and the **Loop** toggle live in the Transport menu without dedicated keys. **A/B Compare Last Edit** (**⌥B**, or **Alt+B** on Windows) flips between your most recent edit and the original of that edit, so you can check the change by ear, then flips back. You can also start, stop, loop, and record from the buttons in the transport bar along the bottom of the window.

File

Action	macOS	Windows
New File	⌘N	Ctrl+N
New from Clipboard	⇧⌘N	Ctrl+Shift+N
Open...	⌘O	Ctrl+O
Close File	⌘W	Ctrl+W
Close All...	⌘⇧W	Ctrl+Alt+W
Show Next File	⇧⌘]	Ctrl+Tab
Show Previous File	⇧⌘[	Ctrl+Shift+Tab
Download Audio...	⇧⌘D	Ctrl+Shift+D
Save	⌘S	Ctrl+S
Save All	⌘⇧S	Ctrl+Alt+S
Save As / Export...	⇧⌘S	Ctrl+Shift+S
Export Selection...	⇧⌘E	Ctrl+Shift+E
Batch Process...	⇧⌘B	Ctrl+Shift+B

Download Audio... fetches audio from a pasted URL — a direct file link, a web page, or a streaming site — and opens it like a dropped file.

Editing

Action	macOS	Windows
Undo	⌘Z	Ctrl+Z
Redo	⇧⌘Z	Ctrl+Shift+Z
Cut	⌘X	Ctrl+X
Copy	⌘C	Ctrl+C
Paste	⌘V	Ctrl+V
Paste ▸ Insert	⌘⇧V	Ctrl+Alt+V
Paste ▸ Replace	⌘⇧⌘V	Ctrl+Alt+Shift+V
Paste ▸ Mix	⇧⌘V	Ctrl+Shift+V
Paste ▸ Invert and Mix	⌘⇧V	Ctrl+Shift+I
Paste ▸ To Selection	⌘⇧⌘V	Ctrl+Shift+T
Delete	Delete	Delete

Paste Special tip: the plain ⌘V (Ctrl+V on Windows) paste behaves the way you last chose; the Paste Special variants let you insert without overwriting, replace a selection, blend the clipboard into what's already there, or fit the clipboard exactly to your current selection.

Selection

Action	macOS	Windows
Select All	⌘A	Ctrl+A
Select None	⇧⌘A	Ctrl+Shift+A
Select Harmonics	⇧⌘H	Ctrl+Shift+H

Select Harmonics takes a tonal selection on the spectrogram and extends it across that note's harmonic series, so you can grab a whistle or hum and everything stacked above it in one move.

In the **Notes** view, **Select All** (⌘A, or Ctrl+A on Windows) selects every detected note in one move rather than the audio buffer — so you can retune, level, or reshape the whole line at once. (Selecting notes is not an audio edit, so it takes no undo step of its own. If a note-name field has focus, ⌘A, or Ctrl+A on Windows selects that field's text instead.)

View, zoom, and tools

Action	macOS	Windows
Waveform view	⌘1	Ctrl+1
Spectrogram view	⌘2	Ctrl+2
Waveform + Spectrogram	⌘3	Ctrl+3
Notes view	⌘4	Ctrl+4
Text view	⌘5	Ctrl+5
Spectrogram Settings	⌘⌥,	Ctrl+Alt+,
Zoom In	⌘=	Ctrl+=
Zoom Out	⌘-	Ctrl+-
Zoom to Fit	⌘0	Ctrl+0
Zoom to Selection	⌘E	Ctrl+E
Amplitude Zoom In	⌘⌥=	Ctrl+Alt+=
Amplitude Zoom Out	⌘⌥-	Ctrl+Alt+-
Time Selection Tool	T	T
Time-Frequency Selection Tool	R	R
Frequency Selection Tool	F	F
Lasso Selection Tool	L	L
Instant Process (toggle)	I	I

The four tool keys (T, R, F, L) and the Instant Process toggle (I) are single letters with no modifier, so they switch instantly while you work over the waveform or spectrogram. Selecting a spectral tool will reveal the spectrogram automatically if it is hidden. (These are the waveform/spectrogram tools; the single-letter tools inside the Notes view are listed separately below.)

Play Single Channel as Dual Mono is a **View** menu toggle (no keyboard equivalent) for multichannel files. It governs how a soloed single channel is monitored — centred on both speakers, or hard-panned to its native output — and is covered under Playback.

Notes editor tools

When the **Notes** view (⌘4, or Ctrl+4 on Windows) has focus, single letters switch the note-editing tool — the keyboard twins of the tool palette in the control strip. They are plain letters with no modifier, so they change tool the instant you press them:

Action	macOS	Windows
Main tool	A	A
Separate tool	S	S
Vibrato tool	V	V
Drift tool	D	D
Intensity tool	I	I
Glide tool	G	G

A letter only switches when that tool is offered for what you're doing: in **Assign** mode only **Main** and **Separate** are available, and **Glide** is mono-only. If a note-name field has focus, the letter types into the field instead of switching tool.

Markers and Regions

Action	macOS	Windows
Add Marker	⇧⌘M	Ctrl+Shift+M
Next Marker	⌘→	Ctrl+Alt+→
Previous Marker	⌘←	Ctrl+Alt+←
Find Similar Events	⌘E	Ctrl+Alt+E
Create Region from Selection	⇧⌘R	Ctrl+Shift+R
Next Region	⇧⌘→	Ctrl+Alt+Shift+→
Previous Region	⇧⌘←	Ctrl+Alt+Shift+←

Find Similar Events scans the file for sounds resembling your current selection and drops markers on the matches.

On the spectrogram

When the spectrogram has focus, a few keys act directly on a spectral selection or repair preview, mirroring the menu commands:

- **Add to / subtract from a selection:** hold $\uparrow$ while you draw to add the new region to what's already selected, or hold ⌘ (Alt on Windows) to carve it out. Without a modifier, the new drawing replaces the selection.
- **Nudge a selection or lasso:** the **arrow keys** move a spectral box, frequency band, or lasso by one pixel; hold $\uparrow$ for a larger 10-pixel step. A held arrow key collapses into a single undo step.
- **Confirm a repair preview:** **Return** (or Enter) applies the previewed spectral repair; **B** auditions the unprocessed signal while previewing; **Esc** cancels the preview (or aborts a drag in progress first); **Delete** removes the lasso point under the pointer.

Module panels

While a restoration or effect module panel is open, the footer's Preview and Bypass buttons carry keyboard equivalents:

Action	macOS	Windows
Preview (render and listen; press again to stop)	$\uparrow$ Space	Shift+Space
Bypass (hear the unprocessed signal while a preview is active)	$\uparrow$ B	Shift+B

Preview is a listen toggle: the first press renders and plays a non-destructive audition of the module; pressed while it is audible it stops the audio but keeps the rendered preview; pressed again it re-listens without re-rendering. **Bypass** A/Bs the processed signal against the original and is active only while an auditionable preview is live. When you are happy, **Process** commits the edit for real (one undoable step). These footer shortcuts now drive previews for every module, not only Spectral Repair.

Gotcha (macOS): plain ⌘ , opens the app's Settings, so Spectrogram Settings is on $\text{⌘}\text{⌘}$, instead, and plain $\text{⌘}M$ is left to the system Minimize, so Add Marker uses $\uparrow\text{⌘}M$. On both platforms, the single-letter tool shortcuts and the module-footer $\uparrow$ Space and $\uparrow$ B only fire when the editing surface (not a text field) has focus.

Windows notes

On Windows the menu bar shows the shortcut it really answers to, so the tables above are the reference. The chords that differ from a straight $\text{⌘} \rightarrow \text{Ctrl}$ translation, and the extras Windows adds:

- **Show Next / Previous File** are **Ctrl+Tab** and **Ctrl+Shift+Tab** (Ctrl+PgDn / Ctrl+PgUp also work) — the macOS $\uparrow\text{⌘}$ / $\uparrow\text{⌘}$ cannot be typed on most European layouts.
- **Paste ▸ Invert and Mix** is **Ctrl+Shift+I** and **Paste ▸ To Selection** is **Ctrl+Shift+T** — $\wedge$ and ⌘ both mean Ctrl on Windows, so the macOS $\wedge\text{⌘}V$ / $\wedge\uparrow\text{⌘}V$ would collide with Paste and Mix.
- **Redo** answers to **Ctrl+Y** as well as Ctrl+Shift+Z.

- **Zoom In / Out** and **Amplitude Zoom In / Out** also answer to the numeric-keypad **Ctrl+Num+** / **Ctrl+Num-** and **Ctrl+Alt+Num+** / **Ctrl+Alt+Num-**.
- **Help ▸ FFT User Manual** is **F1**.
- **Settings** is **File ▸ Preferences...** (**Ctrl+,**); the modifier the tables call ⌘ is the **Alt** key, and **Enter** stands in wherever the tables say Return.

AI Features & On-Device Models

FFT includes a family of AI-powered modules — transcription, stem separation, speech and voice repair, bandwidth regeneration, prompt-based isolation, mastering, text-to-speech, and text-to-audio generation. The key thing to know up front is that **all of it runs on your computer**. There is no cloud round-trip, no account, and no upload of your audio. Every model loads and runs locally, accelerated by your GPU (the Apple Silicon GPU and Neural Engine on macOS; your GPU, or the CPU, on Windows), so your recordings never leave the machine.

How the models work

The AI modules are driven by trained models. A few compact engines are built into the app and ready immediately; the larger ones download on demand — the first time you open a module whose model is missing, its panel shows the model's name, size and license with a download button, live progress, and (where a license requires it) an acknowledgment step before fetching. Once installed, a model stays on your computer and the panel goes straight to its controls.

- **Everything is on-device.** Processing happens entirely on your computer. Downloading a model needs the internet once; after that the AI features keep working with no connection, and nothing about your audio is sent anywhere.
- **Models are local and reused.** Once a model is bundled or installed on your computer it stays available, so subsequent runs of the same module start straight away.
- **Larger models take a moment to warm up.** The bigger restoration and generation engines load into memory the first time you run them in a session. The first pass can feel slower than later ones; that is the model warming up, not a hang.
- **First use shows what's happening.** Where a module is preparing or running, you'll see clear progress (for example a percentage or a progress bar) rather than a frozen window. Long jobs run in the background so the rest of the app stays responsive.

The model library

What each model powers, roughly how big it is, and whether it's built in or fetched on demand:

- **Whisper** — speech recognition for **Transcribe** and the Text editor. Two models download on demand: *English* (~490 MB) and *Multilingual* (~575 MB).
- **Note Detection (Basic Pitch)** — a ~2 MB model bundled with the app. It powers the note detection behind the **Notes editor**, where detected notes can be pitch-edited and exported as MIDI.
- **DeepFilterNet3** — bundled; the fast *Standard* engine in **AI Speech De-noise** (and the Standard AI De-noise stage in **Voice Enhance**).
- **MossFormer2-SE — Studio Voice Restore** (~225 MB, installed on demand) — the *Studio* full-band speech-restoration engine in **AI Speech De-noise** and **Voice Enhance**.
- **MossFormer2-SR — Bandwidth Regeneration** (~440 MB) — the *Studio (MossFormer)* algorithm in **Spectral Recovery**, the Regenerate stage in **Voice Enhance**, and the high-band rebuild in **Resample**.

- **AERO Speech / Music Super-Resolution** (~78 MB each, download on demand) — the *Neural Speech (AERO)* and *Neural Music (AERO)* algorithms in **Spectral Recovery** and **Resample**.
- **LavaSR v2** (installed on demand) — the *Fast (LavaSR)* speech bandwidth-extension algorithm in **Spectral Recovery** and **Resample**.
- **4-stem Separation (HTDemucs)** (~84 MB, installed on demand) — **Stem Remix**'s drums / bass / other / vocals split.
- **Mel-Band RoFormer** — **Voice Isolate**'s vocal/instrumental split, in two variants that download on demand: *Fast* (33 MB) and *HQ* (228 MB).
- **Cocktail-Fork MRX** (~122 MB, download on demand) — **Dialogue Rebalance**'s speech / music / SFX separation.
- **Kokoro, Chatterbox & Qwen3-TTS** — the voices behind **Text to Speech**: the compact Kokoro-82M model for the Preset voices, and the Chatterbox and Qwen studio engines (shared with **Voice Rebuild**) for the natural, multilingual and cloned voices.
- **Stable Audio 3 Medium** (~2 GB) and **ACE-Step XL Turbo** (over 10 GB — by far the largest download) — the two **Generate** engines.
- **Prompt separation engine** — the multi-gigabyte model behind **Prompt Isolate**; its download includes a license-acknowledgment step.

Credits & License

Several AI modules are built on open models from the research community. Each of those panels includes a small **Credits & License** link near the bottom of its controls — for example "SonicMaster (AMAAI Lab) — Credits & License" in **AI Retouch**, and the per-engine credits in **Generate** (Stability AI for Stable Audio 3, ACE Studio for ACE-Step). Click it to open a panel that names the model, the people behind it, and the applicable license, with buttons that open the full license text and the original research paper in your browser. It's there so you can see exactly what powers each feature. The lossy export encoders carry a similar "Encoded on-device · Credits & Licenses" link in the export panel.

Which modules are AI-powered

Each of these has its own chapter with the full controls — this list is a map of where the AI lives:

- **Transcribe** — turns speech into text and subtitles, with a choice of an English-only or a multilingual model.
- **Stem Remix** — separates a mix into vocals, drums, bass and other stems for remixing or export.
- **Voice Isolate** — splits a track into vocal and instrumental stems; keep either, blend them, or export both.
- **Dialogue Rebalance** — separates speech, music and effects, then remixes their levels or exports the stems.
- **AI Speech De-noise** — deep-learning speech denoise, with a fast bundled engine and a stronger full-band restoration engine.
- **Voice Enhance** — a one-control studio voice cleanup chain that uses AI denoise and de-reverb under the hood.
- **Spectral Recovery** — re-synthesizes the missing high band of muffled, phone-quality or band-limited recordings to bring them back to full bandwidth.

- **Prompt Isolate** — describe a sound in words and the AI isolates it or removes it from the mix.
- **AI Retouch** — describe how the take should sound and the AI restores and masters the whole thing (reverb, clipping, EQ, harshness, width) on-device.
- **Distance** — moves a source nearer or further from the mic using AI dereverb and proximity cues.
- **Voice Rebuild** — regenerates a span of speech in the speaker's own voice to fix a damaged word or retake it with new text.
- **Text to Speech** — synthesizes speech from typed text in preset, multilingual, or cloned voices and places it in the timeline.
- **Generate** — creates brand-new audio from a text prompt: insert a clip, inpaint a selection, extend the file, or cover it in a new style.


The **Notes editor** also runs on AI: its note detection (the basis for pitch editing and MIDI export) is powered by the bundled Basic Pitch model.

Tips

- Make a selection first to focus an AI module on just the part you want; with no selection, most of these run on the whole file.
- AI results apply directly to the document and are fully **undoable**, so it's always safe to try a pass and back it out with Undo.
- If a heavy module ever seems slow on the very first run, give it a moment to load — later runs in the same session are quicker.

Settings, License & Updates

FFT's preferences live in one tabbed **Settings** window — choose **FFT ▸ Settings...** (⌘,) on macOS, or **File ▸ Preferences...** (Ctrl+,) on Windows. Three tabs cover app behaviour, the AI model library, and your license and updates.



The General tab

- **Playback** — an **Output device** picker routes playback (the default, **System Default**, follows the system output). Switching re-routes open documents immediately, and a refresh button rescans the hardware. Below it, an **Input device** picker chooses the microphone recording captures from; it takes effect at the next record start, and if the chosen device is missing then, recording falls back to the system default (a caption under the picker reports routing problems). An unplugged device stays listed as *Offline Device* so your choice isn't silently dropped.
- **On Launch** — **Reopen files from the last session** brings back the files that were open when you quit, and **Restore undo history across sessions** lets ⌘Z (Ctrl+Z on Windows) step back through edits made before the app was last closed.
- **Crash Reports** — lists any recent crash logs with a **Reveal** button for each, plus **Clear All Reports**.

The AI Models tab

This tab is the on-demand manager for every AI model and engine the app can use, grouped by what they power — Transcription, Note Detection, Voice Restore & Repair, Stem Separation, and Generation & Mastering — with an installed/total count per group. Each row shows the model's name, its feature, download size and license, plus the action that fits its state:

- **Download** fetches the model (a progress bar tracks it; **Cancel** aborts). Models install into FFT's local data folder (`~/Library/Application Support/Fourier` on macOS, `%LOCALAPPDATA%\Fourier` on Windows) and stay on your computer.
- **Remove** deletes an installed model to reclaim disk space — it simply re-downloads the next time a feature needs it.
- **Bundled** marks models that ship inside the app; a build that bundles everything reports that there is nothing to download.

Some engines built on outside models show a license acknowledgment before their first download. The **Check for Updates** button at the top refreshes the model list itself.

The License & Updates tab

The **License** section shows your activation at a glance:

- **Status** — *Active*, *Active — reconnect within N days* (offline grace), *Needs to reconnect*, *Expired*, *Revoked*, *Seat limit reached*, or *Not activated*.
- Your license key's prefix, the email it's licensed to (when the key carries one), the plan, and how many **machine seats** it allows.
- **Refresh** re-checks the license with the server; **Deactivate This Mac... (Deactivate This Computer... on Windows)** frees this machine's seat so the key can be activated on another computer.

The **Updates** section explains that FFT checks for updates when it opens and tells you when a new version is available — updates are never downloaded or installed automatically. It also holds a **Check for Updates Now** button. The same check lives in the app menu (the **File** menu on Windows) as **Check for Updates...**; it's disabled while a check is already running or the build isn't configured for updates.

Activating FFT

Until the app is licensed, launching it shows a full-window activation screen. The quickest road is **Sign in & Activate**: type the email address you purchased with and click **Email me a code**, enter the 8-digit code from the message, and click **Sign in & Activate** — FFT fetches your license itself and activates this computer, with no key to type. If you'd rather use your key (or sign-in isn't available for your purchase — the screen says so), click **Activate with a license key instead**, enter the key (it looks like `FLUF-XXXXX-XXXXX-XXXXX-XXXXX`) and click **Activate** — name and email are optional, unless your key is bound to an email, in which case the form asks you to add the matching one; **Sign in with your email instead** takes you back. Either way, activation registers this computer and re-checks periodically; the app keeps working offline for a grace period, after which it asks you to connect to the internet and **Retry**.

If all machine seats are already in use, the screen lists each activated computer with when it was last seen — press **Deactivate** next to one to free its seat for this one, or choose **Use a different key**. An expired license asks for a new key or a renewal; a revoked one shows the reason and lets you enter another key.