Hisstory User Guide

A real-time spectral gating plugin designed for removing hiss while keeping high frequency music and transients intact and sounding as natural as possible.

Hisstory was created by Sean Ericson Reach out to Sean with questions, suggestions, or bug reports at **tangotoolkit@gmail.com**

Quick Start

1. Install the Plugin

Windows: Copy `Hisstory.vst3` to `C:\Program Files\Common Files\VST3\`. **MacOS**: Copy `Hisstory.vst3` to `~/Library/Audio/Plug-Ins/VST3/` (or `/Library/Audio/Plug-Ins/VST3/` for all users).

Hisstory includes a **2-day free trial** that starts on first launch. During the trial, all features are fully available — there are no limitations.

- The remaining trial time is shown in the top bar of the plugin
- Click **Get / Enter Key** to purchase a license or enter an existing key
- License keys are 16 characters in the format `XXXX-XXXX-XXXX-XXXX`
- Once a valid key is entered, the trial UI disappears and the plugin is permanently licensed on that machine

When the trial expires without a license key, the plugin passes audio through unprocessed (auto-bypass). Your DAW session and any saved settings are preserved — entering a key restores full functionality immediately.

2. Apply Reduction

- Activate Plugin on Noisy Track
- **Threshold** Determines cutoff for noise, **Reduction** influences amount of reduction.
- Threshold is primary control. Set to remove enough noise without removing too much music.
- Have **Adaptive** ON
- Hear change by using **Bypass** to A/B — keep the result sounding natural

3. Other Useful Controls

- **6-band curve**: drag the numbered points on the analyzer to shape where removal is stronger or weaker
- **Adaptive**: continuously tracks the noise floor as it changes
- **Spectrogram**: visual time-frequency view of what the plugin is doing
- **Advanced** (+): smoothing, dry/wet blend, and residual listen mode
- **?** button: in-plugin help with control and metric descriptions

Tips

- Slightly under-processing almost always sounds more natural than heavy removal
- Right-click any slider to reset it to its default value
- Use Residual Listen periodically as a sanity check — you want to hear only noise in there
- The Dry/Wet control at 85–95% can add back a touch of natural room tone without reintroducing obvious hiss
- For album or playlist processing with Adaptive on, Hisstory re-adapts at each silence gap — no need to reset between tracks
- The spectrogram is especially useful for spotting over-processing: if musical peaks (bright vertical lines) start fading, ease off

Contact

Hisstory was created by Sean Ericson Contact him with any questions, suggestions, or feedback by reaching out at tangotoolkit@gmail.com

How Hisstory Works

1. Split audio into frames Incoming audio is divided into short overlapping windows (4096 samples each, with 75% overlap) and converted into frequency bins using a Fast Fourier Transform.

2. Adaptive noise floor learning When **Adaptive** mode is on, Hisstory continuously builds a per-frequency noise profile during quieter moments. It distinguishes noise from music by measuring how much each frequency bin fluctuates over time — steady energy is noise-like, varying energy is music-like. The learning rate is frequency-dependent: slower in the low-mids (where music lives) and faster in the highs (where hiss dominates). The profile converges to approximately the quietest 15% of energy in each bin, giving a reliable picture of the noise floor.

Hisstory also watches for silence gaps longer than half a second. When one ends, it resets and re-adapts.

3. Spectral subtraction Each frequency bin is compared against the noise floor estimate, offset by the **Threshold** slider and the 6-band curve. Bins that look like noise get attenuated by up to the **Reduction** amount. Hisstory uses soft spectral subtraction rather than a hard on/off gate.

4. Music preservation Hisstory includes several mechanisms to protect musical content

- **Look-ahead transient detection** Hisstory buffers a few frames ahead and scans for sudden energy jumps. When a transient is detected, the gate opens *before* the transient arrives so the attack isn't dulled or suppressed.
- **Tonal peak protection** — Frequency bins that stand out well above their neighbors are identified as tonal content (notes, harmonics, vocal formants). These bins and their harmonic series (2×, 3×, 4× the

fundamental) receive much less attenuation, even if they sit near the noise floor.
- **Broad tonal region detection** — When many adjacent bins are elevated together (chords, bowed strings, sustained vocals), the whole region is protected.
- **Psychoacoustic masking** — Noise that falls below the perceptual masking threshold of a nearby musical tone is inaudible, so Hisstory leaves it alone rather than risk damaging the music to remove something you can't hear anyway. Hisstory uses a Bark-scale spreading function modeled on human hearing.
- **Low-frequency bypass** Everything below 200 Hz passes through ungated. Hiss doesn't live down there, and gating low frequencies can introduce unwanted pumping or rumble artifacts.

5. Smoothing Hisstory applies frequency and temporal smoothing. Multiple smoothing passes across neighboring frequency bins prevent isolated spectral notches or spikes. A gain minimum filter in the 2–12 kHz range catches and fills in narrow dips that would otherwise cause "twinkling" or "birdie" artifacts.

Per-bin gains are smoothed over time with fast attack (gains open quickly) and slower release (gains close gradually). The **Smoothing** slider controls the release speed — higher values give a smoother, less artifact-prone result at the cost of some transient sharpness. Large sudden drops in gain are progressively slowed to prevent modulation artifacts.

6. Overlap-add output The processed frequency bins are converted back to audio and recombined using overlap-add synthesis. The total processing latency is about 7,168 samples (~160 ms at 44.1 kHz).

A final stage ensures the output never exceeds the input level by a significant margin. This prevents amplification artifacts that can sometimes occur during overlap-add reconstruction, especially on transient-heavy material.

Controls

Top Bar

- **Analyzer / Spectrogram** — toggle between two display modes. The analyzer shows real-time frequency curves; the spectrogram shows a time-frequency heatmap
- **Reset** — reset all 6 band control points back to their defaults
- **Adaptive** — continuously learns the noise floor from quiet moments. Recommended for most material. When enabled, Hisstory also detects silence gaps between tracks and re-adapts automatically
- **Bypass** — passes the original audio through unchanged, useful for A/B comparison. Uses a smooth crossfade to avoid clicks when toggling
- **Collapse/Expand** — hide or show the analyzer panel to save screen space

Main Sliders

Threshold [dB] (range: -40 to -10, default: -23)

- Controls how much of the signal gets classified as noise
- Higher (less negative) = more aggressive noise detection
- Lower (more negative) = more conservative, preserves more signal
- You can type values directly into the number box below the slider

Reduction [dB] (range: 0 to 32, default: 12)

- How much detected noise gets attenuated
- Higher values remove more hiss but can introduce dullness or artifacts
- You can type values directly into the number box

6-Band Threshold Curve

The numbered control points on the analyzer let you shape the threshold by frequency region:

Band	Frequency Range
1	200 - 900 Hz
2	900 Hz - 2.1 kHz
3	2.1 - 3.9 kHz
4	3.9 - 6.3 kHz
5	6.3 - 9.8 kHz
6	9.8 - 22 kHz

- Drag a band **up** to remove more noise in that region
- Drag a band **down** to preserve more signal in that region
- This is great for targeting hiss-heavy highs (bands 5-6) while leaving vocals and instruments alone (bands 1-3)

Advanced Settings

Click the **+** button below the metrics to expand the advanced panel:

Smoothing (range: 0 to 100%, default: 50)

- Controls how quickly the noise gate closes after opening
- Higher values give a smoother, more artifact-free sound but can soften transients
- Lower values are more responsive but may introduce modulation artifacts on sustained noise

Dry/Wet (range: 0 to 100%, default: 100)

- Blend between the original (dry) and processed (wet) signal
- At 100%, you hear only the processed output; at 0%, you hear the original
- Mixing in some of the original (e.g. 80–90%) can help preserve natural texture on difficult material where full processing sounds too clean or sterile

Residual Listen

- Lets you hear what's being removed instead of the output
- Ideally you should hear only noise — if you hear music or voice in the residual, back off your settings
- A warning indicator appears when residual listen is active so you don't forget it's on

Help Button

The **?** button opens the built-in help window with tabbed descriptions of all controls and metrics.

Display Modes

Spectrum Analyzer

Shows three real-time curves:

- **Input** (grey) — the incoming spectrum
- **Output** (white) — the processed spectrum
- **Threshold** (orange) — the noise gate curve

Noise reduction happens where the input falls below the threshold line.

Spectrogram

A scrolling time-frequency heatmap using 256 Mel-scale bands. The Mel scale spaces frequencies according to how humans perceive pitch, giving more detail in the lows and mids where most musical content lives. Brighter colors mean louder signal (black → brown → orange → white).

This is useful for checking your work visually:

- Hiss shows up as a persistent colored band across the high frequencies between musical peaks
- Good settings remove that band while keeping the musical peaks intact
- If the peaks themselves are getting dimmer, you're removing too much
- Transient attacks appear as bright vertical lines — these should remain sharp and visible

Metrics

The metrics panel shows real-time indicators with color coding (green = good, orange = caution, red = warning):

HF REMOVED — high-frequency energy (4-16 kHz) removed by processing. More negative means more aggressive reduction. Green: < -3 dB, Orange: -3 to -1 dB, White: > -1 dB.

HARMONIC LOSS — percentage of tonal energy lost above 2 kHz. This is your early warning for musical damage. Green: < 3%, Orange: 3-10%, Red: > 10%.

Troubleshooting

Sound is dull or lispy — lower Threshold (more negative) or lower Reduction. You can also pull down the upper bands (5-6) on the threshold curve. Check the Harmonic Loss metric — if it's orange or red, you're cutting into the music.

Chirpy/warbly artifacts ("musical noise") — back off both Threshold and Reduction a bit. Increasing Smoothing in the advanced panel can also help. These artifacts happen when the gate rapidly opens and closes on individual frequency bins — smoothing tames that.

Not enough hiss removal — raise Threshold slightly (less negative), increase Reduction in small steps, or lift the upper frequency bands on the curve. Make sure Adaptive is on so the plugin has a good read on the noise floor.

Bypass comparison feels jumpy — Hisstory uses a smooth crossfade when toggling bypass, but the level difference between processed and unprocessed can still be jarring. Toggle once and listen to a steady section for a few seconds. Try to level-match by ear for a fair comparison.

Plugin sounds different on a new track — with Adaptive mode on, Hisstory needs a moment to learn the noise profile of new material. It detects silence gaps between tracks and re-adapts automatically, but the first second or two may sound under-processed as the profile builds.

Residual listen shows music content — if you hear recognizable music in the residual (not just noise), your settings are too aggressive. Lower Threshold, lower Reduction, or pull down the band curves in the frequency range where you hear music leaking through.

High latency — Hisstory introduces about 7,168 samples of latency (~160 ms at 44.1 kHz). Most DAWs compensate for this automatically with plugin delay compensation (PDC). If timing feels off, check that your DAW's PDC is enabled.