



Neural**Outboard**

VERSION 3.2 · USER MANUAL

Analog color & saturation, powered by Neural Amp Modeler captures of real hardware.

Contents

1 Overview	2
2 Installation & Activation	3
3 Quick Start	5
4 The Interface	6
5 NAM Sets	13
6 Presets, Your Own Content & Expansion Packs	14
7 Updates	16
8 Troubleshooting	17
9 Specifications	18
10 Credits & Legal	19

1 Overview

NeuralOutboard is an analog-coloring and saturation processor for mixing and mastering. Instead of modelling one fixed circuit, it loads NAM captures — neural models of real outboard hardware (preamps, compressors, tape machines, EQs, console channels) trained with Neural Amp Modeler. Drop a capture in, drive it to taste, and you get that unit's harmonic character, compression and EQ on whatever you run through it: a full mix, a drum bus, or a single track.

Read the panel top-to-bottom and the signal flow follows — trim the input and pick the routing, drive the loaded model(s), oversample to keep the saturation clean, rebalance mid/side, blend dry and wet, and set the final output, with DC-blocking and full input/output metering along the way.

Input trim > Routing > Drive → Model > Oversample > Mid/Side > Dry/Wet > Output



The NeuralOutboard interface, with the hardware view open and a capture loaded in both slots.

2 Installation & Activation

Installing — macOS

Open the downloaded disk image and run the installer. It installs both plug-in formats system-wide, plus the manual and licence documents:

Item	Location
Audio Unit	/Library/Audio/Plug-Ins/Components/NeuralOutboard.component
VST3	/Library/Audio/Plug-Ins/VST3/NeuralOutboard.vst3
Documents	/Library/Application Support/dbdsp audio/NeuralOutboard/

NeuralOutboard is a universal binary and runs natively on Apple Silicon and Intel Macs.

Installing — Windows

Run the downloaded installer. Windows hosts use VST3, so that is the only plug-in format installed:

Item	Location
VST3	C:\Program Files\Common Files\VST3\NeuralOutboard.vst3
Documents	C:\Program Files\dbdsp audio\NeuralOutboard\

The manual, licence agreement and third-party notices also appear in the Start Menu under **dbdsp audio**. Windows 10 or later, 64-bit.

After installing, restart your DAW (or rescan plug-ins) so it picks up the new plug-in.

Trial & activation

Open the preset menu (click the preset field) and choose **License...** to open the license panel. From there:

- **Start a free 14-day trial** — one press, no key needed. The trial is the complete, unrestricted plug-in (one computer per trial). When it lapses, the plug-in locks until a license key is entered.
- **Activate** — paste the license key from your purchase receipt and press Activate. A purchased license covers up to three computers at a time.
- **Manage Licenses** — once activated, the panel's left button becomes Manage Licenses and opens your customer portal, where you can view orders, download the latest version, and deactivate machines.
- **Moving to a new machine** — deactivate an old computer (via Manage Licenses) to free an activation, then activate on the new one.
- **Reset this machine** — a local escape hatch (tap twice to confirm) that clears this computer's activation state, e.g. after a machine was deactivated remotely. It does not consume or release seats

by itself.

Internet. Activation and periodic license re-verification need an internet connection; if the plug-in can't reach the license server, check your connection and firewall settings. Once activated, the plug-in tolerates being offline for a reasonable grace period and keeps working. See section 9 for exactly what the plug-in does — and doesn't — send.

3 Quick Start

Five steps to a colored, level-matched sound:

- Load a capture into Slot A (and Slot B for stereo or mid-side work) — or pick a preset from an installed pack.
- Press **Auto Input** to land in the ballpark, then ride Input and Drive to taste for the source.
- Press **Auto Gain** so your A/B is loudness-matched, then toggle **Bypass** to judge tone honestly.
- Use **Delta** to hear exactly what the plug-in is adding or removing.
- Want more grit? Reach for a more colorful capture — the knobs can only deliver what the model was captured with.

New in 3.2 — GPU acceleration. Two buttons in the header run the neural model on your graphics card instead of the CPU: **GPU** for mixing, which adds one buffer of latency and buys the most headroom, and **GPU +0ms** for tracking, which adds no latency at all. See section 4.10.

NAM Sets. A single preset can hold a whole series of captures that you sweep between with one Morph knob — e.g. a drive rig captured at several input levels, morphed smoothly from clean to slammed. Load one from a pack's Factory NAM Sets menu, or build your own. See section 5.

4 The Interface

4.1 Top bar

Presets · Info · < > · Save · Save As

The preset field shows the current patch; click it to open the full preset menu (section 6). The < > arrows step through presets. **Save** updates the current preset; **Save As** stores the full plug-in state — knobs, routing, oversampling and loaded captures — as a new preset. **Info** opens the about panel. A preset references its captures by file, so keep your .nam files in place.

The SAVE button flashes whenever the loaded preset has unsaved edits. Audition controls — Bypass and Delta — deliberately don't count as edits, and neither do the GPU buttons; everything else does, including a committed Auto Input or Auto Gain result.

Bypass

True bypass — passes the untouched dry signal for comparison. The plug-in reports its latency to the host, so bypass stays time-aligned, and the switch itself is a short crossfade (about 5 ms), so toggling is click-free even on loud program material. Bypass always has the last word: engaging it while Delta is active gives you the clean dry signal.

4.2 Session helpers

Auto Input · one-press · target –21.1 dBFS RMS

One-press input calibration. It measures the broadband RMS of your incoming signal over the analysis window and sets the Input trim so the signal sits at the reference level NAM models are characterised at (–21.1 dBFS RMS). While it measures, the trims preview live; the result commits when the window completes — press the button a second time to lock the measurement early. Treat it as a “get me in the ballpark” starting point, not a fixed truth.

Window knob · 3–30 s

The small knob beside the button sets the analysis window. A longer window averages more material for a steadier reading on dynamic sources.

Delta

Outputs only the difference between the processed and dry signals — precisely what the plug-in is adding or removing (harmonics, EQ shift, compression artefacts). The fastest way to hear what your coloration is actually doing.

HQ · full model vs. light model

Red (on, default): the capture runs at its full network size — best quality. Off: the model runs slimmed, trading a little fidelity for a lighter CPU load — useful when tracking with many instances. Toggling produces a deliberate, brief (≈8 ms) dip while the engine switches network size — that's the click-free handover, not a glitch. HQ is the plug-in face of the automatable Slim parameter; the engine treats it as a two-state switch, so automation flips cleanly at the halfway point.

Auto Gain · K-weighted (LUFS) match

One-press output level-matching. It compares the K-weighted loudness (ITU-R BS.1770 / LUFS) of the input versus the processed output and sets Output so the two match in perceived loudness — an honest, level-matched A/B: toggle Bypass and you compare tone, not volume.

4.3 Gain & saturation

These are the controls you'll live on — on captured gear, input level is the instrument, exactly as on the real hardware.

Input L / Input R · primary drive · link

Trim the level going into the model, per channel. Because a NAM model is nonlinear, input level sets how hard you drive the captured gear — and therefore how much saturation and compression you get, within the range the gear was captured at. The chain icon links L and R so they move as one (the normal choice for stereo material); unlink to balance channels independently.

M-S Pre · Stereo & Mid-Side modes

A mid/side balance applied before the model: tilt the drive between the mid (centre) and side (stereo-difference) components — for example, push the sides harder than the centre.

Drive · global · “–” = unity

The red macro — a global push into the model(s) on top of the Input trims. Up hits the gear harder for more saturation and compression; down stays clean. Remember the ceiling: a capture can only deliver the saturation baked into it.

2× Drive · small red tab, top-left of Drive

Off (default) keeps Drive's range at ± 24 dB. On, it doubles the effective drive range to ± 48 dB, for captures that want to be pushed much harder (or backed off much further) than the standard range allows. Reach for it only when the normal range runs out of room; most sources never need it.

Morph · NAM Sets only · small knob below Drive

The small knob just below and right of Drive sweeps through the members of a loaded NAM Set — for example, a rig captured at several drive levels, morphed continuously from the cleanest member to the most saturated. It has no effect until a set is loaded. Below the knob sits a small round set-offload toggle: red means the whole set is resident and the Morph knob is live; gray collapses the set to just the single model at the current Morph position, freeing CPU and memory. Full detail in section 5.

M-S Post · Stereo & Mid-Side modes

Mid/side balance applied after the model: rebalance centre versus sides in the processed signal. Pairs with M-S Pre for full mid-side control of the colour. The chain icon between the M-S Pre and M-S Post knobs links them as mirror images — move one and the other tracks inversely, for quick symmetric mid/side tilts.

D/W Mix · 0–100% · parallel blend

The dry/wet blend, set after the model and the mid/side post stage. Pull it back from 100% to mix the clean signal in underneath for parallel saturation — keeping dry transients and clarity while the wet path adds harmonic density and glue. The blend works with every anti-alias filter at every oversampling factor.

Output · makeup gain

Final makeup gain at the end of the chain. Match the processed level back to your dry level (or let Auto Gain set it). Saturation usually adds level, so this typically trims down.

4.4 Hardware view

The panel between the knobs and the capture slots shows a photo of the hardware behind each loaded capture — Slot A on the left, Slot B on the right. The small chevron at the far left of the knob row collapses or reveals the view (the window resizes with it), and the choice is remembered.

- Expansion packs carry their images built in — pack presets show their hardware automatically.
- Your own captures show a photo when an image file with the same name sits beside the .nam file — see section 6.2 for the exact convention.

4.5 Capture slots

Two independent capture slots. Click a slot to browse and load a .nam capture; the small square button beside it clears the slot. Each loaded slot shows an architecture chip (the model's network type, e.g. A2, A1, LSTM) and, when the capture carries loudness metadata, its output level — a reliable readout of how hot the capture runs, handy for judging makeup gain.

Slot labels follow the processing mode: Slot A reads Mono / Left / Mid and Slot B reads Right / Side. In Stereo, A processes Left and B Right; in Mid-Side, A processes Mid and B Side — so you can run different gear on centre versus stereo content. Load the same capture in both for uniform processing.

Model swap — AA / AB / BA / BB · under the Slot B field

A four-way selector beneath the Slot B field decides which loaded model feeds which processing channel, without you having to reload anything. AB (default) is the natural mapping. BA swaps them; AA runs model A on both channels; BB runs model B on both. Handy for auditioning your two captures against each other on the same content.

4.6 Channel routing

Left / Right / Mid / Side · solo · monitoring

Solo a single component to audition it in isolation — solo Side to hear only what's happening to your stereo content, Mid to focus on the centre. A monitoring aid; it doesn't change the printed output.

Mono · Stereo · Mid-Side

Mono: single-path processing. Stereo: dual-mono — model A on Left, model B on Right. Mid-Side: encodes to Mid and Side, processes each independently (A = Mid, B = Side), then decodes back — saturate the centre differently from the stereo ambience.

4.7 Oversampling

Saturation creates harmonics above the audio band; without oversampling they fold back as aliasing — harsh, inharmonic artefacts. Oversampling runs the nonlinear stage at a higher internal rate so those harmonics are filtered out cleanly before downsampling.

Real Time vs Offline · two independent settings

Real Time applies during playback and tracking — balance quality against CPU. Offline applies during bounce/render, so you can print with a higher, more expensive factor without bogging down live playback. Each pane's status bar reads back the active factor and filter, plus the DC Block corner and its added latency for that engine.

Link · chain between the two panes

The chain icon between the Real Time and Offline panes links the two sections (on by default) so the oversampling factor, filter, and DC Block settings you dial for playback are mirrored to the render engine automatically — set it once and print matches what you hear. Unlink whenever you want a heavier configuration for the final bounce than you run live.

Factor · Off / 2× / 4× / 8× / 16× / 32×

How many times the capture's internal rate is multiplied. Higher means less aliasing but more CPU. 2×–4× is plenty for gentle colour; 8× and up for heavy drive. If your machine has a capable graphics card, section 4.10 makes the higher factors far cheaper in real time.

What “Off” really means. At OS Off the model runs at its native rate — for nearly all NAM captures, 48 kHz. If your session runs at the same rate, the signal passes straight through with no conversion and no filter in the path. If your session runs at a different rate (44.1, 88.2, 96 kHz...), the plug-in converts sample rate around the model, and your selected anti-alias filter governs that conversion — so the filter choice audibly matters even at “Off”.

Filter · IIR / Min FIR / Poly / Linear

The anti-alias filter used at the rate conversion. IIR (minimum-phase) and Min FIR are low-latency with minimal pre-ringing — good for tracking. Poly (polyphase FIR) and Linear (linear-phase) are phase-coherent and steeper, at the cost of latency and possible pre-ringing — better for mixing and mastering.

4.8 DC Block

A clean-up high-pass that removes DC offset and the subsonic build-up saturation can introduce. The strip below the oversampling section has two selectors — the left one for the Real Time engine, the right one for the Offline (bounce) engine — sharing one cutoff knob:

- **Off** bypasses the filter for that engine.
- **Min** (minimum-phase) is a gentle, low-latency high-pass that keeps transient timing intact — the everyday choice.
- **Linear** (linear-phase) is phase-coherent with no low-end smearing, at the cost of latency and possible pre-ringing — best for mastering or mid-side work.

The cutoff knob sets the corner frequency (2–200 Hz; you can also click a readout and type a value). Keep it low to strip only DC and subsonics; raise it to clear unwanted low end. Each oversampling pane's status bar shows the corner and that engine's added filter latency (e.g. 2 Hz · 0.0 ms).

4.9 Metering

Each bar overlays average and peak level; the two numbers read RMS first, then peak (e.g. $-11.4 / -7.6$). The gap between them is the crest factor — transient versus body — which lets you set input drive and output makeup without clipping. Click a meter to reset its held peak.

4.10 GPU acceleration

Two buttons in the header hand the neural model to your graphics card instead of the CPU. On a machine with a capable GPU this is dramatically cheaper — enough to change what's practical at high oversampling, where a single instance can saturate a CPU core.

GPU

The model runs on the graphics card and the plug-in reports one extra buffer of latency to your host. Delay compensation handles it, so nothing drifts out of time; you simply have one buffer more round-trip delay than the CPU path. This is the mode for mixing and mastering, where a buffer of latency costs you nothing and the extra headroom lets you stack instances or run heavy oversampling that would otherwise be out of reach.

GPU +0ms

The same processing with no added latency at all. The plug-in waits for the graphics card within the same processing block rather than returning the previous one. That costs capacity — you get considerably fewer instances than the GPU button gives — but it adds nothing to your monitoring path. This is the mode for tracking and live monitoring, and it still runs comfortably more instances than the CPU can at the same zero latency.

The two are mutually exclusive: clicking one turns the other off, and clicking the lit one returns to CPU. The change takes effect immediately.

The setting follows the machine, not the preset. Whether a model runs on the GPU is a property of the computer you're working on, not of the sound. So changing presets, loading a factory pack or opening a NAM Set never changes it, and a preset you share never forces GPU mode on someone whose machine can't support it. The plug-in remembers your last choice and opens new instances that way, while a saved project restores whatever that project was using.

If the graphics card stops responding

A driver reset, a device removed, some other fault — the plug-in switches every instance back to the CPU, tells you it has done so, and carries on. You lose the offload, not the audio.

Is it the same sound?

Yes. The GPU path is verified against the CPU path to the limit of 32-bit floating point — a difference far below anything audible, and far below the difference between two takes of the same performance.

When it won't help

Laptops with graphics built into the processor share memory bandwidth with the CPU, so offloading may gain little. If a model can't be moved to the graphics card at all, the plug-in tells you why and keeps running on the CPU. Requirements are in section 9.

4.11 Bottom bar

The bottom line shows the running version and format (e.g. v3.2.0 (VST3)). When a newer version of NeuralOutboard is available, a yellow tab appears over the wordmark — click it to open the download page. See section 7.

5 NAM Sets

A NAM Set is a single preset that contains a series of related captures instead of just one — and one Morph knob to sweep between them. Think of a drive pedal captured at five input settings, or a compressor captured at several ratios: a set bundles those snapshots into one slot so you can move smoothly across the range with a single control (and automate it) rather than swapping presets.

5.1 Loading and morphing a set

Sets load like any other content. A factory pack's sets live in the **Factory NAM Sets** menu; your own live in **User NAM Sets** (section 6). Load one and the slot behaves as a single capture, with two additions:

- The **Morph** knob (below Drive) sweeps continuously through the set's members. At the far left you're on the first member; at the far right, the last; in between, the plug-in blends across them so the transition is smooth rather than stepped.
- The **set-offload toggle** (the small round button under the Morph knob) controls how much of the set is held in memory. Red — the full set is resident and Morph is live. Gray — the set collapses to just the single model at the current Morph position, releasing the CPU and memory the rest of the set was using.

Morph is automatable. Because the sweep is one continuous parameter, you can automate a move across the set from your DAW — ride from a clean capture into a driven one across a section, for instance. Keep the set expanded (toggle red) wherever the automation is moving.

5.2 Loading a set from a folder

The preset menu's **Load NAM Set...** item builds a set on the fly from a folder of captures — point it at a folder whose .nam files are the members you want to morph between, and it loads them into the slot as a set. This is the quick way to audition a series of captures as a set without packaging anything.

5.3 Saving your own sets

When a set is loaded and you press **Save As**, NeuralOutboard stores it as a user set — a single .no3set file — and it appears under **User NAM Sets** in the menu, kept separate from your ordinary single-capture presets.

6 Presets, Your Own Content & Expansion Packs

6.1 The preset menu

Click the preset field to open the menu. From top to bottom:

Item	What it does
Default	The clean starting state.
Factory Packs	One submenu per installed pack of single-capture presets, organised by category. Locked packs show a lock glyph — see 6.3.
Factory NAM Sets	Installed packs whose presets are NAM Sets (section 5), kept in their own menu so sets and single presets don't mix.
User Presets	The single-capture presets you've saved with Save As.
User NAM Sets	The NAM Sets you've saved (your .no3set files).
Delete current	Removes the loaded user preset or set.
Reveal folder	Opens your preset folder in Finder or File Explorer.
Load NAM Set...	Builds a set from a folder of captures on the fly (5.2).
User Manual	Opens this manual.
Update available...	Appears only when a newer version exists.
License...	Opens the license panel (trial, activation, Manage Licenses).

Your presets are stored at `~/Library/Application Support/NeuralOutboard/Presets/` on macOS, and `%LOCALAPPDATA%\NeuralOutboard\Presets\` on Windows — that is `C:\Users\<you>\AppData\Local\NeuralOutboard\Presets\`. A preset references its captures by file path — if you move or delete a .nam, presets that use it can't load it. Keep your capture library in a stable location.

6.2 Your own presets — with hardware photos

Any .nam capture you load can show its hardware in the photo view. The convention is one simple rule: **same folder, same name**. Place an image next to the capture with the identical file name and a .png, .jpg or .jpeg extension. `MyPreamp.nam + MyPreamp.png` → the photo appears whenever that capture is loaded.

That's the whole workflow for personal content: keep each capture beside its photo, load it into a slot, dial the sound, and Save As. The preset remembers the capture; the capture's photo follows it automatically — including when the preset is recalled in a saved session.

6.3 Expansion packs

Packs are collections of captures — with photos and categories built in — installed by their own installer: a .pkg inside a disk image on macOS, an .exe on Windows. Installed packs appear in the preset menu automatically; no restart of the plug-in is needed beyond reopening its window. A pack of single captures lands under Factory Packs; a pack of sets under Factory NAM Sets.

Unlocking a paid pack

A paid pack shows a lock glyph in the menu. Selecting one of its presets opens the unlock panel: paste the pack's key from your receipt and press Unlock — the preset you clicked loads the moment the key validates. Unlocking needs an internet connection. An unlocked pack stays unlocked — including across reinstalls.

Pack trials & upgrades

Packs may offer time-limited trial keys. A trial unlocks the full pack until it lapses, after which the lock returns; entering a purchased key at any time — including over an expired trial — unlocks the pack permanently.

Removing packs

Each pack's installer package includes the dbdsp Preset Packs uninstaller: run it, tick the packs to remove, and it takes care of the rest. Your unlocks are kept by default, so reinstalling a pack later doesn't ask for the key again.

7 Updates

NeuralOutboard tells you when a newer version exists — quietly. At most once per day it fetches a small version file from dbdspaudio.com; when the published version is newer than the one you're running, two indicators appear: a yellow tab over the wordmark at the bottom of the plug-in, and an “Update available — v...” entry in the preset menu. Click either to open the download page. Nothing updates itself; you stay in control.

Installing an update

Download the new version and run its installer — it installs directly over the old one. Your license activation, pack unlocks, user presets, and installed packs are all untouched by updates, and the update notice clears itself once the new version is running. Some hosts revalidate plug-ins after an update; if the new version doesn't appear, rescan plug-ins or restart the DAW.

Privacy. The update check is a plain web request. Like any web request it discloses your IP address and a generic software identifier to the web server — nothing else: no account, usage, audio, or personal data. Offline, the check fails silently and the plug-in simply carries on. Full details in the Privacy Policy at dbdspaudio.com.

8 Troubleshooting

Symptom	What to check
Plug-in doesn't appear in the DAW	Rescan plug-ins or restart the host. In Logic, run a plug-in rescan; in other hosts check both the AU and VST3 lists (VST3 only on Windows).
"Demo Expired" panel	The 14-day trial has lapsed. Enter a purchased key to activate, or use Manage Licenses if you already own one.
Key won't activate — "no seats left"	All three activations are in use. Open Manage Licenses and deactivate a machine you no longer use, then try again.
"Can't reach the license server"	Activation and pack unlocks need internet. Check your connection and firewall settings — some studio machines block plug-in network access.
Pack shows locked again after its trial	That's the trial lapsing as designed. Enter the purchased key to unlock permanently.
Morph knob does nothing	Morph only acts on a loaded NAM Set. If a set is loaded, make sure the set-offload toggle is red (expanded) — gray collapses the set and freezes the sweep.
A saved set landed in the wrong menu	Sets you save appear under User NAM Sets, single presets under User Presets — the plug-in sorts by what the slot holds when you Save As.
No hardware photo for my capture	The image must sit in the same folder as the .nam with the identical name and a .png / .jpg / .jpeg extension (section 6.2).
Preset loads but a slot is empty	Presets reference captures by file path — the .nam was moved or deleted. Restore it or re-load and re-save.
GPU button lights but nothing sounds different	It should not sound different — GPU processing is verified identical to the CPU path. What changes is how much you can run.
GPU button won't stay on	The plug-in switches back to the CPU when a model can't be moved to the graphics card, and tells you why. Integrated graphics may simply not be worth using.
Extra latency after enabling GPU	Expected — the GPU button adds one buffer, compensated by your host. Use GPU +0ms if you need none.
Brief level dip when toggling HQ	Intentional — the engine ducks for ≈8 ms to switch model size without a click.
Meters hold an old peak	Click the meter to reset its peak hold.

Anything else: info@dbdspaudio.com — include your operating system and version, DAW and version, and what you were doing.

9 Specifications

Plug-in name	NeuralOutboard — analog color & saturation processor
Version	3.2.0
Manufacturer	dbdsp audio · dbdspaudio.com
Formats	macOS: Audio Unit (AU) · VST3 — Windows: VST3
Platform	macOS · universal binary (Apple Silicon + Intel) — Windows 10 or later, 64-bit
Engine	Neural Amp Modeler inference core with oversampled nonlinear processing; selectable full/slim network size (HQ)
GPU acceleration	Optional, per instance. Two modes: standard (one buffer of added latency, maximum capacity) and zero-latency. Automatic fallback to the CPU if the graphics card cannot be used. Verified against the CPU path to 32-bit float precision. Windows: requires a graphics card supporting Direct3D 12. macOS: requires Metal, which every supported Mac provides.
Capture format	.nam (Neural Amp Modeler); optional same-name .png / .jpg / .jpeg hardware photo
NAM Sets	Multi-capture presets with a continuous Morph sweep and CPU-saving set-offload; user sets saved as .no3set
Reference level	−21.1 dBFS RMS (NAM model characterisation level)
Processing modes	Mono · Stereo · Mid-Side · AA/AB/BA/BB model swap
Drive range	±24 dB, or ±48 dB with 2× Drive engaged
Oversampling	Independent (optionally linked) Real-Time and Offline engines · Off–32× · IIR / Min FIR / Poly / Linear filters
DC Block	Independent Real-Time and Offline modes · Off / Min / Linear · 2–200 Hz shared cutoff
Latency	Measured per configuration and reported to the host. GPU mode adds exactly one buffer; GPU +0ms adds none.
Licensing	Free 14-day full-feature trial (one computer) · purchased license covers up to 3 computers · sold via Lemon Squeezy
Network use	License validation/re-verification · once-daily update check (no telemetry; see Privacy Policy)
Expansion content	Installable preset and NAM-Set packs with built-in hardware photos; per-pack unlock keys and optional trials

10 Credits & Legal

NeuralOutboard is built on Neural Amp Modeler by Steven Atkinson (model format and inference core). Oversampling follows the method developed by DLC86 (NAM-Oversampler). Third-party open-source components and their licenses are listed in the “Third-Party Software Notices and Licenses” file installed alongside the End User License Agreement that governs your use of the software — in /Library/Application Support/dbdsp audio/NeuralOutboard/ on macOS, and C:\Program Files\dbdsp audio\NeuralOutboard\ on Windows.

VST is a trademark of Steinberg Media Technologies GmbH, registered in Europe and other countries. Audio Units, macOS, Metal and Apple are trademarks of Apple Inc. Direct3D, DirectX and Windows are trademarks of Microsoft Corporation. Neural Amp Modeler (NAM) is an open-source project created by Steven Atkinson; dbdsp audio and NeuralOutboard are not affiliated with, sponsored by, or endorsed by the Neural Amp Modeler project, Steinberg, Microsoft, or Apple.

© 2026 dbdsp audio · NeuralOutboard v3.2.0 · This document may be redistributed with the plug-in.