



NEURALDREAM

SYNTHESIZER, NEURAL AND DSP
INSTRUMENT AND FX PROCESSOR
USER MANUAL (V2.17)

BEFORE YOU BEGIN...

THANK YOU FOR PURCHASING/TRIALING NEURALDREAM!

Before You Begin: Trial and Activation

The first time you open NeuralDream, a licence panel headed **Welcome to NeuralDream** covers the plugin window. Unless NeuralDream is on a trial or activated, it loads in your DAW but it makes no sound. One licence covers both NeuralDream Instrument and NeuralDream FX.

Path 1: Try NeuralDream free for 14 days

To start a trial, choose **Try it free for 14 days**, type your email address and press **CONTINUE**. NeuralDream then emails you a 6-digit code. The panel changes to Check your email. Type/Paste in the code and press **CONTINUE**. The code lasts for 15 minutes. If it does not arrive, look in your junk or spam folder. You can also press **Send a new code** for a fresh one.

When the trial has started, the panel shows the date it runs until. Press **CONTINUE TO NEURALDREAM** to start working. During the trial, a notice beside the version number (at the bottom right of the plugin's window) shows how many days are left. It turns amber for the last three days and red on the final day. The free trial can be started once on each computer.

When the trial ends, the licence panel returns and NeuralDream will be silent again. If you would like to Activate then please use the **BUY A LICENCE AT DREAMDEEPCR.COM**.

Path 2: Activate with your licence key

When you buy NeuralDream you receive a licence key. To activate, choose **I have a licence**, type the email address you bought NeuralDream with and the key and press **CONTINUE**. The panel then confirms that NeuralDream is activated and names the email address it is licenced to. You can activate at any time, whether or not you started a trial first. During a trial, click the trial notice at any time to activate NeuralDream with your purchased licence key.

If the email address or the key is not right, the panel says they do not match, so check both against your purchase email. Each licence key can be activated on up to three computers.

Path 3: A Beta licence file from DreamDeep

We sometimes issue a licence for one computer directly, for example to beta testers. To receive one, click the **Beta tester?** link at the bottom right of the licence panel, click **COPY COMPUTER ID** and send the copied ID to DreamDeep. When the licence file arrives, click **LOAD LICENCE FILE** on the same page and choose it. NeuralDream checks that the file was issued by DreamDeep for this computer before it uses it and starts working straight away. A licence file made for another computer is refused.

A licence of this kind can have an end date. NeuralDream stops at the end of that day, even while your DAW is open. The panel then shows the date your licence ended. To carry on, type your purchased licence key and press **CONTINUE**.

Other important information

An internet connection is needed only twice

Starting the trial and activation needs an internet connection once. Loading a licence file from DreamDeep needs none. After that, NeuralDream never needs to be online: it checks its licence from a file on your own computer when it opens.

Keep your computer's clock correct

NeuralDream remembers the latest date and time it has seen. If your computer's clock is set back earlier than that, the panel asks you to set the date and time correctly and reopen the plugin and NeuralDream stays silent until you do.

Files you should not delete

Your licence is stored in two files, **license.dd** and **hwm.dd**. On a Mac they are in *~/Library/Application Support/DreamDeep/NeuralDream* and on Windows they are in *%APPDATA%\DreamDeep\NeuralDream*. Do not delete, move or edit either file.

The file **license.dd** is your licence itself; if it is deleted you will have to activate NeuralDream again and a trial cannot be started a second time. **hwm.dd** works alongside it to protect the licence.

NeuralDream also keeps your general preferences, such as the voice limit and the background choice, in a small settings file called **NeuralDream.settings**. On a Mac it is in *~/Library/Application Support/DreamDeep* and on Windows it is in *%APPDATA%\DreamDeep*. Do not delete it either or those preferences will return to their defaults. The instrument and the effect keep their own preferences inside that one file, so you can give each of them its own backdrop, its own voice limit and its own panel and visualiser choices even when both are on the same track. The one preference they share is the author name used when you save a preset, because that describes you rather than a plugin.

1. INTRODUCTION

1.1 WHAT IS NEURALDREAM?

NeuralDream is a modular audio processing plugin for your DAW that reimagines how musicians shape sound. Rather than offering a single fixed effect, it provides a multi-layered architecture where each layer allows you to create entire DSP and Neural FX chains of *Audio Transformers* and even link layers together.

Audio Transformers are not just *effects*, although some operate like this (with just one audio source and a panel whose controls can be fully modulated from a variety of sources) but, many require two audio sources, we call these carriers and modulators. Your carrier audio can be drawn from the plugin's input, its sidechain, samples or the output of other layers and modifiers can be samples, impulse responses, downloaded trained ML latent spaces such as accurate guitar amp profiles, etc.


With 40 (or so) engines spanning filtering, delay, reverb, granular processing, binaural processing, spectral manipulation (multiple), spectral freezing, vocoding, convolution, neural network effects, ring modulation, pitch shifting (with scales) and more; NeuralDream gives composers and producers an open-ended sound design environment where complex processing chains can be built, saved and recalled as presets - all within a single plugin instance.

What sets NeuralDream apart is its deep integration of machine learning directly into the creative signal path. Four neural engines - Neural FX, Neural Amp, Machine Learnt Tape Simulation and RAVE use real-time inference to process audio in ways that conventional DSP simply cannot replicate. Neural FX generates randomised temporal convolutional networks from a seed number, meaning every seed is a unique effect. Neural Amp loads .nam model files trained on real amplifier circuits using the Neural Amp Modeller framework. RAVE uses variational autoencoders to transform audio through learned latent spaces, enabling timbral transformations that blur the line between two sound sources.

Modulation brings everything to life. Four LFOs, four envelope followers and four macro controls can be routed to any parameter across any engine, bringing static processing chains to life with movement and expression. A 1024-slot parameter pool ensures that every knob in every engine is a potential modulation target and a potential automation lane in your DAW. Sample playback is a first-class feature with six playback modes, DAW transport sync, time stretching with phase-locked vocoder processing, automatic BPM detection and an interactive waveform display. In the Instrument, a built-in step sequencer plays looping gate, pitch and velocity patterns - plus four extra modulation lanes - driving the synthesiser exactly as if the notes had been played over MIDI.

NeuralDream is built for the modern composer and producer who wants to move beyond preset-surfing and into genuine sonic exploration. Its machine learning engines do not replace musical intuition - they amplify it, offering starting points that no amount of conventional knob-twisting could reach, while keeping every parameter under direct creative control.

NeuralDream FX and **NeuralDream Instrument** is designed by **DreamDeep** and built with JUCE.



IR Convolver using a supplied Impulse Response (IR)

1.1.1 WHAT IS NEURALDREAM INSTRUMENT?

NeuralDream Instrument is a companion plugin that shares the same engine architecture, modulation system, preset format and user interface as NeuralDream but operates as an extraordinarily powerful polyphonic software instrument rather than an audio effect.

Where the effect version processes audio arriving from your DAW, the instrument version generates its own audio from loaded samples triggered by MIDI notes, making it a fully playable synthesiser using very advanced Wavetable, Grains, Partial and Modal based synthesis, with the complete NeuralDream processing chain available on every layer.

NeuralDream Instrument is available as an AU Software Instrument on macOS and as a VST3 Instrument on Windows. Both versions are built from the same source and share all features.

1.2 INSTALLATION

1.2.1 MACOS

NeuralDream ships as a macOS installer (.pkg) inside a DMG disk image. Double-click the DMG to mount it, download this manual from the DMG and double-click the installer package. The installer offers four optional components, each independently selectable:

Audio Unit (AU) Installed to /Library/Audio/Plug-Ins/Components/ for use in Logic Pro, GarageBand and other AU-compatible hosts.

VST3 Installed to /Library/Audio/Plug-Ins/VST3/ for use in Ableton Live, Reaper, Cubase, FL Studio and other VST3-compatible hosts.

Factory Presets Installed to /Library/Application Support/DreamDeep/NeuralDream/ and available to all users on the Mac. These are read-only reference presets that demonstrate the range of the plugin.

Impulse Responses A curated library of impulse response WAV files installed to /Library/Application Support/DreamDeep/NeuralDream/Impulses/. These include room, hall, plate, spring and creative impulse responses for use with the IR Convolver engine, plus two sets of vintage speaker cabinet IRs.

Starling subject files (**.star**) - saved subjects and their settings, stored alongside your presets and sequences in the DreamDeep/NeuralDream folder. See Section 11.8.5.

User presets that you create are saved to ~/Library/Application Support/DreamDeep/NeuralDream/ and coexist alongside the factory presets. If a user preset has the same filename as a factory preset, the user version takes priority - this allows you to customise factory presets without modifying the originals.

1.2.2 WINDOWS

NeuralDream for Windows ships as a zip file holding the installer, this manual and the licence text. Unzip it, double-click the Setup file and follow its steps, clicking Yes if Windows asks whether to allow it to make changes to your device. The installer puts the following in place:

VST3 Both plugins, NeuralDream and NeuralDream Instrument, installed to C:\Program Files\Common Files\VST3\ for use in Ableton Live, Reaper, Cubase, FL Studio and other VST3-compatible hosts.

Factory Presets Installed to C:\ProgramData\DreamDeep\NeuralDream\ and available to all users on the PC. These are read-only reference presets that demonstrate the range of the plugin.

WaveTables The factory wavetable collection, installed to C:\ProgramData\DreamDeep\NeuralDream\WaveTables\.

Impulse Responses The curated impulse response library for the IR Convolver engine, installed to C:\ProgramData\DreamDeep\NeuralDream\Impulses\. Windows hides the ProgramData folder, so to browse to it from the IR Convolver's LOAD button, type the path into the address bar at the top of the window.

Manual and Licence This manual and the licence text, installed to C:\Program Files\DreamDeep\NeuralDream\.

User presets that you create are saved to C:\Users\<your name>\AppData\Roaming\DreamDeep\NeuralDream\. When the installer has finished, open your DAW and rescan its plugins.

1.2.3 NEURALDREAM INSTRUMENT INSTALLATION

NeuralDream Instrument installs alongside the effect version. On macOS the installer places both the AU effect and AU instrument in their standard locations and on Windows it places both VST3 plugins in C:\Program Files\Common Files\VST3. Both versions share the same factory presets, impulse responses and user preset directory. Presets are tagged as "effect" or "instrument" and the preset browser filters automatically to show only presets compatible with the version you are running.

1.2.4 WHERE NEURALDREAM LIVES ON DISK

Beyond the plugin binaries themselves, NeuralDream keeps its content in two places: a machine-wide folder that the installer creates and never changes afterwards and a per-user folder that the plugin creates and grows as you work. Knowing which is which tells you exactly what is safe to move, back up or delete.

Mac

Installed by the package (machine-wide, read-only - **/Library/Application Support/DreamDeep/NeuralDream/**):

- **Preset files (.ndpreset)** and **Sequence files (.seq)** – the factory preset library and a small number of factory sequences, available to every user on the Mac.
- **WaveTables** – the factory wavetable collection, one folder per pack.
- **Impulses** – the curated impulse-response library for the IR Convolver.

Created and grown as you work (per-user - **~/Library/Application Support/DreamDeep/NeuralDream/**):

- **Preset and Sequencer file** – presets and sequences you save. A user preset sharing a factory preset's filename takes priority over it, so you can customise factory sounds without touching the originals.
- **WaveTables** – your own wavetable additions. Anything you place here appears in the wavetable browser and a user file with the same relative path as a factory file overrides it.
- **WaveTables/Imported** – created on demand: audio you import lands here. Because presets embed the imported tables they use, this folder is a self-healing cache rather than a store - you can delete it at any time and any preset that needs one of its files quietly recreates that file on load.
- **WaveTableCache (.ndmip)** – pre-computed playback tables **These are large by design:** each wavetable is baked into a full set of anti-aliased frames for every playing register, so a single long import can occupy several hundred megabytes here. It is purely derived data - **the WaveTableCache folder can be deleted at any time to reclaim disk space** and NeuralDream simply rebuilds each table the next time it is loaded, at the cost of a few seconds' wait per table.

Windows

On Windows the same layout lives under C:\ProgramData\DreamDeep\NeuralDream (machine-wide) and C:\Users\<your name>\AppData\Roaming\DreamDeep\NeuralDream (per-user), with the same rules throughout.

1.3 NEURAL NETWORK MODELS

NeuralDream integrates four neural network engines that use real-time inference to process audio. Three of these - Neural Amp, the neural mode of Tape Simulation and RAVE - require external model files that are not included in the installer due to their size and third-party licensing. The fourth, Neural FX, generates its own networks procedurally from a seed number and requires no external files.

NAM MODELS (.NAM FILES)

Used by the Neural Amp and Neural+DSP Tape Sim engines. NAM (Neural Amp Modeller) is an open-source project by Steven Atkinson that uses deep learning to capture the complete nonlinear behaviour of analogue equipment - amplifiers, pedals, preamps, tape machines and more. The trained model files are compact (typically 50-200KB) and run in real time with negligible latency. You can download hundreds of free community-trained models from:

NAM supports multiple neural network architectures - WaveNet, LSTM, ConvNet and Linear - and NeuralDream handles all of them transparently. Models vary in CPU cost: WaveNet models are the most accurate but heaviest while Linear models are lightweight and suitable for less complex tonal captures.

You can download hundreds of free community-created NAM models from ToneHunt (<https://tonehunt.org>) and TONE3000 (<https://www.tone3000.com>).

To load a NAM model click the LOAD button on the Neural Amp or Tape Simulation panel. The model name is displayed on the panel and saved with your preset so the correct model is recalled automatically when you reload the preset.

RAVE MODELS (.TS FILES)

Used by the RAVE engine. RAVE (Realtime Audio Variational autoEncoder) is a deep learning model developed by Antoine Caillon and Philippe Esling at ACIDS, IRCAM, Paris (published 2021). Unlike conventional effects that apply fixed transformations RAVE learns the statistical structure of a body of audio - speech, instruments, environmental sounds, noise - and uses that learned representation to transform whatever audio you feed it in real time.

The result is unlike anything achievable with traditional DSP: audio is encoded into a compact latent space (typically 8 dimensions) where it can be manipulated and then decoded back into audio that carries the tonal and timbral character of the training data. Feed a guitar through a RAVE model trained on violin recordings and the output takes on string-like qualities while following the dynamics and phrasing of the guitar performance.

Pre-trained RAVE models are available from ACIDS IRCAM (https://acids-ircam.github.io/rave_models_download). Download a .ts (TorchScript) model file and load it via the LOAD button on the RAVE panel. Note that RAVE requires the ML-enabled build of NeuralDream (built with LibTorch).

1.4 SYSTEM REQUIREMENTS

macOS: macOS: Apple Silicon (M1 or later) running macOS Big Sur (11) or later. NeuralDream is built as a native ARM64 binary - no Rosetta translation required.

Windows: NeuralDream is available as a VST3 plugin for 64-bit Windows 10 or later.

CPU: NeuralDream's CPU usage varies significantly depending on which engines are active. Simple effects like the Low Pass Filter or Noise Gate are very lightweight. FFT-based engines (Spectral Morph, Vocoder, Pitch Shifter) and neural engines (Neural FX, Neural Amp, RAVE) are more demanding. Multiple layers with complex chains will increase CPU load proportionally. In the synthesiser, load rises with the number of notes you hold and a rich patch - three oscillators with wavefolding, both filters in use, heavy modulation - costs considerably more per note than a simple one.

***Note:** If you hear crackling, crunching or dropouts as you play additional notes, increase the audio buffer size in your DAW's audio settings before assuming anything is wrong. Moving from, say, 256 to 512 samples doubles the time available to render each block and will usually clear it completely. NeuralDream already spreads its voices across multiple processor cores but the DAW still has to receive each block by a fixed deadline and a small buffer makes that deadline too short sometimes for all of the maths that is going on inside the plugin (which is already highly optimised).*

RAM: NeuralDream itself uses modest memory. However loaded samples (up to 10 minutes of audio per sample slot), impulse responses (up to 30 seconds) and neural network models (particularly RAVE .ts files) will increase memory usage. A minimum of 8 GB RAM is recommended.

DAW Compatibility: NeuralDream works with any DAW that supports Audio Unit (AU) or VST3 plugins including Logic Pro, Ableton Live, Cubase, Studio One, Reaper, Bitwig and others.

1.5 ABOUT SOME OF THE TECHNOLOGY

NeuralDream's engines draw on decades of audio signal processing research and several recent advances in machine learning for audio. Here are some of the key techniques and their origins.

Phase Vocoder and Pitch Shifting - The pitch shifter and time stretch engine use the STFT (Short-Time Fourier Transform) framework with identity phase locking as described by Laroche and Dolson (1999) in "Improved Phase Vocoder Time-Scale Modification of Audio" published in IEEE Transactions on Speech and Audio Processing. Transient detection follows Roebel (2003) in "A New Approach to Transient Detection" from the Proceedings of DAFX, using spectral flux with adaptive smoothing to preserve the attack quality of percussive sounds during time stretching. The core pitch shifting algorithm is based on Bernsee's smbPitchShift (WOL licence). Formant preservation (only invoked if Formant is in a non-zero.default position) uses Roebel-Rodet cepstral envelope estimation per-frame. Steps: (a) compute log-magnitudes, (b) inverse FFT → real cepstrum, (c) zero all quefrency bins above cutoff (low-pass in cepstrum domain = smooth envelope), (d) forward FFT → envelope in linear-magnitude space (via exp), (e) flatten source by dividing anaMagn by envelope, (f) shift the now-flat spectrum as today, (g) multiply shifted bins by the envelope, optionally stretched. Formant control in -100...+100%, 0 = bypass (zero overhead). Cutoff: 30 quefrency bins, floor: 1e-10. Reference: Roebel & Rodet, DAFX 2005.

Tape Simulation - The hysteresis model uses the Jiles-Atherton equations solved with fourth-order Runge-Kutta (RK4) integration at 4x oversampling following the approach described in the DAFX textbook. The complete tape path includes hysteresis, head bump, high-frequency loss filtering, wow and flutter modulation and tape hiss - each stage modelled independently and combined in series. In neural mode a NAM model replaces the hysteresis stage while the remaining analogue tape characteristics (loss, bump, wow/flutter, noise) continue to run as conventional DSP around it.

FDN Reverb - The feedback delay network uses a Hadamard mixing matrix following Jot and Chaigne (1991) in "Digital Delay Networks for Designing Artificial Reverberators" from the Proceedings of the AES Convention. Eight mutually prime delay lines with per-line damping produce dense, natural-sounding reverberation.

The tail is set as a true decay time, with each line given the feedback that reaches it; the mixing can be varied from the full Hadamard matrix down to eight separate lines by a network of rotations that keeps the energy of the sound intact at every setting; and the shimmer uses the same phase-locked pitch shifter as Lustre.

Noise Gate - Implements the feedforward gate topology described by Giannoulis, Massberg and Reiss (2012) in "Digital Dynamic Range Compressor Design - A Tutorial and Analysis" published in the Journal of the Audio Engineering Society Vol. 60 No. 6. The implementation features peak detection, 3 dB hysteresis to prevent chattering at the threshold boundary, configurable hold time and one-pole attack/release smoothing with coefficients derived directly from the paper's equations.

Spectral Style Transfer - Cross-synthesis via spectral envelope transfer following J.O. Smith III's work at CCRMA, Stanford (2020). The algorithm extracts spectral envelopes from carrier and modulator signals using frequency-domain running-average smoothing of log-magnitudes then transfers the modulator's spectral shape onto the carrier while preserving the carrier's phase structure.

VCF - The voltage-controlled filter uses a topology-preserving transform (TPT) state variable filter with per-sample smoothing in the log-frequency domain following the approach described by Vadim Zavalishin in "The Art of VA Filter Design" and inspired by the Huovilainen ladder filter model.

Chorus - Based on the bucket-brigade device (BBD) delay line model with modulated delay times producing the characteristic ensemble thickening effect of analog chorus units.

Neural Amp Modeller - Uses the official NeuralAmpModelerCore library by Steven Atkinson (MIT licence) supporting WaveNet, LSTM, ConvNet and Linear architectures for real-time amp, cabinet and pedal emulation.

RAVE - Realtime Audio Variational autoEncoder by Caillon and Esling (2021) from the ACIDS lab at IRCAM, Paris. A deep generative model that encodes audio into a low-dimensional latent space and decodes it back, enabling real-time timbral transformation. Published as "RAVE: A variational autoencoder for fast and high-quality neural audio synthesis."

Neural FX (RONN) - A randomised temporal convolutional network that generates unique nonlinear distortions, saturations and textural transformations from a single seed number. Each seed produces a deterministic network with different learned characteristics giving access to an essentially infinite palette of neural effects without requiring any external model files.

1.6 ALIASING AND NEURALDREAM

A digital system can only represent frequencies up to half the sample rate. If a process creates energy above that line it does not vanish and it is not quietly filtered away: it folds back down and reappears at a frequency that is the mirror image of where it should have been. The trouble is that the reflected tone bears no musical relationship to the note you played. Harmonic distortion adds tones that are whole multiples of the fundamental and that is why it sounds like warmth or drive; aliasing adds tones at arbitrary intervals below the note and they move the wrong way when you play up the keyboard. That is the giveaway. Play a slow rising scale on a bright patch and listen for something descending underneath it or for a metallic shimmer that gets worse the higher you go.

There are two places it comes from. Any waveform with sharp corners, a sawtooth or a square, has harmonics going up forever; play one at 5 kHz and everything above its tenth harmonic folds back into the audible range. And any non-linear process, distortion, wavefolding, hard clipping, bit reduction, ring modulation, manufactures new harmonics out of whatever it is given, so it can alias even when what it was given was perfectly clean.

There are correspondingly two cures:

1.6.1 BAND-LIMITING

Builds the waveform so that the offending harmonics are never created in the first place: either by applying a small correction at each discontinuity or by holding a set of pre-filtered tables, one per octave or part of an octave and playing the one whose content is safe for the note being sounded.

1.6.2 OVERSAMPLING

Takes the other approach: it runs the process at two or four times the sample rate, so that the folding point moves far above anything you can hear and then filters carefully and returns to the normal rate. Band-limiting is the tool for oscillators, where the harmonic content is known in advance. Oversampling is the tool for non-linear processes, where it is not.

1.6.3 NEURALDREAM INSTRUMENT AND EFFECTS

NeuralDream uses both of the above techniques, in the places each belongs.

Every wavetable is stored as a band-limited pyramid of twenty-four half-octave bands per frame, built once and shared, so the oscillator always plays a table whose content is already safe for the pitch it is sounding. The Modal engine's exciter oscillators are corrected at each discontinuity by the PolyBLEP method, which is why its sawtooth and square stay clean at the top of their range.

On the non-linear side, the Tape Simulator runs its hysteresis model at four times the sample rate, the Grit processor runs at four times and the synthesiser's VCF filter model runs at twice, each with a proper filter on the way back down. The Exciter's harmonic generators run at four times the sample rate or twice at 88.2 kHz and above and are built so that they can never create anything above three times the frequency they are given, so nothing they create can alias.

Being straightforward about the rest: not every non-linear process in the instrument is oversampled and it would be dishonest to claim otherwise. The saturation stages inside several of the effect engines and the deliberate destruction in Bitfold, run at the host rate. In Bitfold's case that is the point; sample-rate reduction and bit crushing are aliasing, used musically and removing it would remove the effect. In the others the non-linearity is gentle enough that what it produces above the folding point is far below the signal. The line we draw is this: where a process is meant to sound clean, it is protected by band-limiting or oversampling and where a process is meant to sound broken, it is left alone on purpose.

2. PLUGIN OVERVIEW

2.1 THE INTERFACE AT A GLANCE

NeuralDream's interface is organised as a vertical stack of functional zones. From top to bottom you will find the visualiser, preset strip, the layer area with engine panels, the modulation strip and finally the black 'trim' strip, which contains the master input and output trims, meters and master kill switch indicator (triggered at +24dB, you are allowed to create positive feedback loops in NeuralDream!).

Each zone has a distinct role and together they give you full control over every aspect of the plugin's processing without needing to navigate between pages or tabs.

Every knob in NeuralDream can be set precisely with the Shift key. Hold Shift while you drag a knob and it moves at a quarter of its normal speed, however fast or slowly you move the mouse. You can press or let go of Shift at any point during a drag and the knob carries on smoothly from where it is, without jumping.

2.2 THE VISUALISER

The visualiser occupies a topmost area and provides real-time visual feedback on the audio flowing through the Transformer that has the current focus (clicked on):



Three view buttons control what is displayed. IN and OUT appear only in the effect version of NeuralDream; the instrument offers XFM alone:

IN - Shows the input waveform arriving at the plugin.

OUT - Shows the final output waveform after all processing.

XFM - Shows the engine-specific visualisation for whichever transformer is currently selected. Each of the 40+ engines has its own unique visualiser - the DARKSKY engine shows a 3D aurora, the Tape Simulation engine shows a hysteresis curve, the Parametric EQ shows a frequency response plot with coloured bands, the Noise Gate shows a scrolling level history with gate state and so on. A layer whose source is set to Particles shows its own cloud here: a three-dimensional box in which every living particle is a point of light. Switch XFM off (and IN and OUT in the effect version) and the visualiser folds away to its row of buttons, giving its space to the layers below - useful on a smaller screen.

Where a display shows the pitch or waveform of a single note - the synthesiser's wave display and its frequency readout in particular - it follows the note you played most recently. This matters on a resonant patch, where an earlier note may still be ringing louder than the one you have just struck: the display switches to the new note regardless, so the readout always describes what you are playing rather than whatever happens to be loudest at that instant.


[] - or double-click anywhere on the visualiser to toggle fullscreen mode. In fullscreen the visualiser expands to fill the entire plugin window, hiding the layer area and modulation strip. This is useful for performance or for getting a detailed view of an engine's behaviour. Use the left and right arrow keys to step through presets while in fullscreen mode. Double-click again to return to the normal view.

HIDE | SHOW - Hide or show the modulation strip with its LFOs, macros and envelope followers to give your layers more room or just to enjoy the cosmic nebula background! NeuralDream remembers your choice and the strip opens the same way next time unless a preset has its own setting for it.

There is a unique visualiser for every audio transformer in NeuralDream.

2.3 THE PRESET STRIP

The preset strip runs across the top of the plugin and provides everything you need to manage presets.



From left to right it contains:

INIT - Resets the entire plugin to its default state. All layers are removed except one, all engines are cleared and all modulation routings are disconnected. This is your clean slate.

SCAN - Refreshes the preset list by rescanning the factory and user preset directories. Use this if you have added new preset files while the plugin is running.

R (Randomise) - Loads a random sub-preset into every active engine in the patch. The engines themselves stay where they are; only their settings change and only to settings that already exist as sub-presets. It is a quick way to hear what the current combination of engines can do. Undo restores what you had.

V (Variant) - Generates settings you have not tried, which is the difference between V and R: R shuffles between existing sub-presets, V invents new values for the controls of the engine you are working on. Click V to vary the engine that currently has the focus. Right-click V to vary every engine in the patch at once. The synthesiser is treated as an engine too, so a variant of the instrument reshapes oscillators, filters, envelopes and modulation together. Every press stores an undo state first, so any variant is one Undo away and the button is safe to press freely. If you have unsaved work the first press asks whether to continue; once you have said yes, further presses in the same session do not ask again. V and its AMOUNT knob appear only in the instrument version of NeuralDream. The effect plugin does not show them and leaves their space in the strip empty. V works on the synthesiser, so in the instrument it is greyed out when no layer has a Synth source. SEQUENCER and STARLING are greyed in the same circumstances, which includes the whole of the effect plugin. Hovering either of them says why.

AMOUNT (for V randomness) - The small knob beside V sets how far a variant may stray from the current settings, from 1 to 100 percent, with 25 as the starting point. At low values a variant is a nudge that keeps the character of the patch; at high values it can land somewhere else entirely. Some controls are deliberately left alone or held within a safe range by V, so a variant never produces a broken patch: a reverb is never locked into freeze, a delay is never pushed into runaway feedback and a wavetable oscillator is always given a table to play.

Preset Navigation - The left and right arrows step through presets one at a time. The preset name is displayed between the arrows. Presets are organised by category and you can browse the full library by clicking the preset name to open the preset browser.

SAVE - Saves the current state as a user preset. The preset browser opens in save mode, where you give the preset a name, an author and a description of up to 1000 characters, click one of the six categories and click as many tags as apply. The description appears in the browser's preview strip whenever the preset is selected and is searchable. See Section 3.3.

Undo/Redo - The arrow buttons to the right of SAVE step backward and forward through your recent changes. NeuralDream maintains a 12-level undo history covering parameter changes, engine additions, engine removals and preset loads.

BYPASS - Bypasses all processing across the entire plugin, passing audio through unchanged.

CLR - A safety-guarded clear button. Press and hold for 1/2 second to trigger a full reset. A progress bar fills as you hold and the button flashes green when the clear completes. This clears all state including stopping any sample playback.

2.3.1 'V' - THE VARIANT GENERATOR

The V button generates settings you have not tried. It is the difference between V and R: R shuffles between sub-presets that already exist, so the engines stay where they are and only their settings change and only to settings that someone has saved. V invents new values for the controls of an engine, within limits that keep the result musical. Every press stores an undo state first, so any variant is one Undo away and the button is safe to press freely. If you have unsaved work the first press asks whether to continue; once you have said yes, further presses in the same session do not ask again. V and its AMOUNT knob appear only in the instrument version of NeuralDream.

Which engine it varies. Click V to vary the engine that currently has the focus, which is the one whose panel is highlighted. Right-click V to vary every engine in the patch at once. The synthesiser is treated as an engine too, so a variant of the instrument reshapes its oscillators, filters, envelopes and modulation together and a variant of the whole patch reshapes everything.

How far it goes. The small AMOUNT knob beside V sets how far a variant may stray from the current settings, from 1 to 100 percent, with 25 as the starting point. At low values a variant is a nudge that keeps the character of the patch: the same sound with different proportions. At high values it can land somewhere else entirely. A good way to work is to press V at a low amount several times, listening and raise the amount when the results stop surprising you.

What it leaves alone. Not every control is a fair target for chance and V follows a table of rules that say, for each control of each engine, what it may do. Most controls are varied freely across their range. Some are bounded to a safe part of it: every delay's feedback is held below fifty percent, so a variant never produces runaway echo and a filter cutoff is kept where it still lets sound through. Some are excluded altogether, because changing them would break the patch rather than vary it: every engine's Freeze switch, a compressor's makeup gain, the Resonance Field's PLAY SYNTH and POLY switches and the synthesiser's master output level are never touched. A few are varied only occasionally, because they change the character so much that a variant with them always on would sound like the same instrument every time: portamento on the synthesiser arrives one press in five and stays short, the filter link one press in four, the Vowel filter's voice and depth one press in five. And a wavetable oscillator is always given a wavetable to play, so a variant never produces a silent oscillator.


Modulation. On the synthesiser a variant may also create modulation routings, not only change knob positions: an LFO to a filter cutoff, an envelope follower to a level, chosen from the controls that are the most musical things to modulate, with the filter cutoffs weighted to be chosen most often. So a variant of the instrument can move as well as sound different.

What V is for. It is a way of exploring a sound rather than designing it: press, listen, undo or keep, press again. Because it works within the rules above, the results are playable rather than random and because it works per engine, you can vary the reverb while leaving a sound you like intact or vary the synthesiser while keeping the effects you have set up. Save anything you want to keep as a preset; V does not save for you.

2.3.2 PRESET MANAGER

The preset browser opens when you click the preset name in the preset strip. It fills most of the plugin window and is where you find, audition, describe and save presets.

Across the top is a band carrying an image for the category you are looking at. Choose a category and the pictures at the top and the foot change to it at once, before you touch a preset. With All chosen or the favourites star, they follow whichever preset you are pointing at instead. The search field sits on that band. Type into it and the list narrows from the third character onward, matching the preset name, its description, its author and its tags, so typing "wheel" finds the preset whose description mentions the mod wheel. Beside the search field, SORT orders the list by name A to Z, name Z to A, newest first or oldest first, which is what you want after an evening of making presets.



Below the band are four columns:

LIBRARY narrows the list by the kind of sound source a preset uses, with six buttons in two rows: ALL, SAMPLER and SYNTH on the top row and PARTICLES, SIDECHAIN and NOISE below. Clicking one shows only the presets containing a layer with that source.

CATEGORY filters by the six categories or by your favourites.

ENGINES lists every engine in the plugin, grouped exactly as the add-engine menu groups them, so you can ask for every preset that uses RAVE.

TAGS holds the whole tag vocabulary in four groups; clicking tags narrows the list to presets carrying all of them, so PAD and MOTION together means both.

All four choices are remembered while the plugin is open, so a search for a kind of sound does not have to be set up again every time you open the browser. Close the browser, work on your sound and open it again and your library, category, engines and tags are still set as you left them. The previous and next preset arrows beside the preset name step through exactly the presets the browser is showing, so they follow the same filters, tags included. Nothing is remembered between sessions: closing your DAW clears it all.

CLEAR, the button centred under the tags, empties the filters in one go: the category goes back to All and every engine and every tag is switched off. Use it when a preset you have just saved does not appear in the list, which usually means a filter you set earlier is hiding it. Along the foot is the preview strip, which describes whichever preset you are pointing at: its name, its category, its author, its tags and its full description. The same category image runs behind it, mirrored, so the panel is bracketed by colour.

A single click loads a preset and leaves the browser open, so you can work through several in succession. A double click loads it and closes the browser. The up and down arrow keys move through the list, loading each preset as you land on it and Return loads the highlighted preset and closes the browser. Each preset row also carries a star to make it a favourite and, for your own presets, a bin to delete it.


Factory presets and user presets appear together. If a user preset shares a filename with a factory preset then only your version is shown, which lets you keep personalised variations of factory sounds without losing the originals.

2.4 THE LAYER AREA

Below the visualiser sits the Layer Area where your processing chains are built. Each layer appears as a horizontal strip showing a safety-controlled 'x' button to delete the layer, the source selector on the right, the transformer chain in the centre, input level and input meter, Solo and Mute buttons and, finally, the output level / pan controls and output meter. You can insert up to 8 engines in any layer which means the layers also side scroll and show a scroll bar when necessary. You can have up to 8 layers too.

The source selector lets you choose where a layer draws its audio from - Plugin Input (the audio arriving from your DAW), Sidechain, Sample (a loaded audio file), Noise (a built-in coloured noise generator) or the output of another layer.

The transformer chain shows each engine as a compact coloured block. Click any engine to select it and its control panel appears below the layer strip. You can add new engines in a layer by clicking the + button at the end of the chain. Engines can be dragged to infer a reorder using the handle icon on each engine's header. use the little 'x' button to delete a layer. To the right of each layer strip you will find the MIX knob (dry/wet balance for the entire layer), the LEVEL knob and the Solo, Mute and Bypass buttons. LEVEL is the layer's output fader: most of its travel is a ± 24 dB trim around unity but the very bottom of the travel is true silence - the knob's readout in the Helper switches to OFF at that point and the layer's output (and its meter) goes genuinely to



zero. The transition into and out of silence is gain-ramped, so riding the fader to nothing never clicks. Solo isolates a single layer while muting all others; Mute silences the layer's contribution to the output; Bypass passes the source audio through unchanged, skipping all chained engines.

New layers are added using the '+ ADD LAYER' button below the layer area. Right-click a layer header for additional options including deletion. The first layer cannot be deleted.

2.5 AUDIO TRANSFORMER (ENGINE) PANELS

When you select an engine in the layer area its control panel appears directly below. Every engine panel follows a consistent layout - a header bar showing the engine name and sub-preset controls, followed by rows of knobs specific to that engine.


Each engine has its own accent colour and many feature an overlay image that gives the panel a distinctive visual identity. Knobs show their current value in the readout area when dragged. All knob values can be automated by your DAW and modulated by NeuralDream's LFOs, envelope followers and macros.

Sub-presets let you save and recall just the settings for a single engine independently of the full plugin preset. This is useful when you want to reuse a particular delay configuration or EQ curve across different presets.

To the right of each audio transformer panel, you will see a combination of mix and level controls. MIX adjusts the output between pure dry input and full wet engine output. LEVEL is a post-engine trim and the BYPASS button (B) is smooth-ramped per-engine bypass that includes bypassing MIX and LEVEL values. Some engines need a moment to do their work and hand their sound back slightly later than they received it. MIX always blends the engine's output with a copy of its input that has been held back by exactly the same amount, so blending never causes the hollow, phasey sound of two copies playing out of step. Bypass passes on that same held-back copy, so switching an engine in and out never shifts the timing of its layer. 2.6 The Modulation Strip

The modulation strip runs across the bottom of the plugin and contains three blocks:


LFO Block - Four independent LFOs (Low Frequency Oscillators) with rate, depth and waveform controls. Each LFO can run in free mode at a set frequency or sync to your DAW's tempo with beat divisions ranging from 1/32 to 4 bars. Available waveforms include sine, triangle, sawtooth, square and sample-and-hold. LFOs are sample-accurate when synced to transport - they lock precisely to your DAW's playback position.



Macro Block - Four macro knobs that act as master controllers. Each macro can be routed to multiple parameters simultaneously so a single knob turn can sweep a filter cutoff, increase a delay feedback and reduce a reverb mix all at once. Macros are ideal for live performance and for creating expressive preset variations.



Envelope Follower Block - Four independent envelope followers that track the amplitude of the audio signal and convert it into a control signal. Use these to make parameters respond dynamically to the music - for example driving a filter cutoff from the input level so the filter opens on loud passages and closes on quiet ones.



COPY **PASTE** **FILTER 1 CUTOFF**

Synth LFO 1

+50%

UNI Q P 1 OCT VIA x

EF 1

29%

UNI Q P 1 OCT VIA x

VIA SLOW

Macro 4

63%

UNI Q P 1 OCT VIA x

+ Add source...

Finally, the MOD button in the top right of the panel allows you to close (and hide when clicked again) the Modulation Area, to have more Layer access to enjoy the Nebulous Space background!

2.7.1 HOW IT WORKS

Each layer maintains its own voice allocation. Different layers can load different samples with independent polyphony settings. Every layer hears every MIDI note so you can stack multiple sample layers to create rich layered instrument sounds, each processed through its own chain of transformer engines.

The instrument version adds no new engines. All 40 (or so) transformer engines from the effect version are available and work identically. The difference is purely in how audio enters the processing chain: via MIDI-triggered sample playback rather than a live audio input.

2.7.2 THE EXTENDED SAMPLE CONTROL PANEL

When a layer's source is set to Sample in the instrument version the sample control panel expands to show additional controls not present in the effect version. The panel grows from four columns to nine columns to accommodate the amplitude envelope, filter section, velocity control and voice count.

The left section (columns 0 to 3) contains the same playback controls as the effect version: START, END, SPEED, LP START, LP END, XFADE, PLAYBACK, MODE, STRETCH, SHIFT and FFT. These work identically to the effect version and are described in Chapter 5.

A vertical divider separates the playback section from the instrument controls.

2.7.2.1 AMPLITUDE ENVELOPE (ADSR 1)

Four knobs control the per-voice amplitude envelope that shapes the volume of each note from the moment it is triggered to the moment it fades to silence.

AMP ATK (Attack) -- How quickly the note reaches full volume after a key is pressed, from 0.1 ms to 10 seconds. Short attack times (under 10 ms) give immediate, percussive onsets. Longer times create swelling, pad-like entries.

AMP DEC (Decay) -- How quickly the volume drops from the peak to the sustain level, from 0.1 ms to 10 seconds.

AMP SUS (Sustain) -- The volume level maintained while a key is held down, from 0% to 100%. At 100% the note stays at full volume for as long as the key is held. At lower values the note decays to a quieter sustained level.

AMP REL (Release) -- How quickly the note fades to silence after the key is released, from 0.1 ms to 10 seconds. Short release times cut the note cleanly. Longer times create a natural fade-out tail.

The envelope uses exponential curves with a slight overshoot on the attack stage for a natural analog-style shape. A voice is freed and becomes available for reuse only when its amplitude envelope reaches silence after the release stage completes.

2.7.2.2 FILTER SECTION

A per-voice filter with its own ADSR envelope provides timbral shaping that responds to each note independently.

CUTOFF -- Filter cutoff frequency from 20 Hz to 20 kHz.

RESO (Resonance) -- Emphasis at the cutoff frequency from 0% to 100%.

DEPTH -- Filter envelope modulation depth (bipolar). Positive values cause the filter envelope to open the filter above the cutoff setting. Negative values cause it to close below the cutoff. At 0% the filter envelope has no effect and the cutoff remains static.

KEY TRK (Key Tracking) -- Controls how much the filter cutoff follows the MIDI note pitch, from 0% to 100%. At 100% the cutoff tracks the keyboard one-to-one so higher notes have a brighter timbre. At 0% the cutoff is the same for every note.

Filter type buttons select between four modes: LP (low pass), HP (high pass), BP (band pass) and COMB (comb filter).

2.7.2.3 FILTER ENVELOPE (ADSR 2)

FLT ATK, FLT DEC, FLT SUS, FLT REL -- These four knobs control a second independent ADSR envelope that modulates the filter cutoff according to the DEPTH setting. The filter envelope has the same range and behaviour as the amplitude envelope but drives the filter rather than the volume.

A common technique is to set a short filter attack with moderate decay and low sustain so that each note starts bright and quickly darkens, then use the DEPTH knob to control how dramatic this effect is.

2.7.2.4 VELOCITY

VEL -- A bipolar velocity curve control. At the centre position (0%) velocity response is linear: soft key presses produce proportionally softer notes. Turning the knob left (negative values) compresses the velocity curve so soft notes play closer in volume to loud notes. Turning it right (positive values) expands the curve so soft notes become much quieter while hard notes remain at full volume. The formula uses a power curve: $\text{gain} = \text{velocity raised to the power of 4 raised to the curve amount}$. Velocity 127 always produces full volume and velocity 0 always produces silence regardless of the curve setting.

2.7.2.5 VOICE COUNT

The VOICES button cycles through four polyphony settings: Mono (1 voice), 8, 16 and Max (32 voices). In Mono mode each new note cuts the previous note. In polyphonic modes new notes are always allocated to a fresh voice and old voices playing the same MIDI note enter their release stage rather than being cut. When all voices are in use the oldest voice is stolen.

Voices render asynchronously on separate CPU cores using Grand Central Dispatch (macOS) for efficient multi-core utilisation. This pipelined architecture introduces one audio block of latency (approximately 5.8 ms at 256 samples / 44.1 kHz) which is imperceptible in practice.

2.7.2.6 ROOT NOTE DETECTION

When a sample is loaded NeuralDream attempts to detect the root note from the filename. Common naming conventions are recognised: note names (C, C#, Db, D, etc.) optionally followed by an octave number (0 to 8) and optionally followed by key quality indicators (m, min, minor, maj, major). Examples: "Piano_C3", "Pad_F#4", "Bass_Ebm", "Strings_Aminor2".

The convention follows Logic Pro's standard where C3 equals MIDI note 60 (middle C). If no root note is detected the default is C3.

The detected root note determines the pitch relationship between MIDI notes and sample playback. Playing the root note plays the sample at its original pitch. Notes above the root play faster (higher pitch) and notes below play slower (lower pitch) unless time stretching is enabled.

2.7.2.7 TIME STRETCHING IN INSTRUMENT MODE

When the STRETCH button is enabled in instrument mode the behaviour differs from the effect version. In effect mode stretch synchronises sample playback speed to the DAW tempo using the ratio between the sample's original BPM and the current DAW BPM. In instrument mode there is no BPM tempo matching. Instead stretch enables pitch-independent playback: every MIDI note plays the sample at its original duration regardless of pitch.

Without stretch (the default) the instrument behaves like a classic sampler: pitch changes speed. Playing an octave above the root note plays the sample at double speed and half duration. This is the vinyl record or tape speed-change behaviour familiar from hardware samplers.

With stretch enabled a two-stage phase vocoder processes the audio. The first stage time-stretches to compensate for the pitch change. The second stage resamples the output to achieve the target pitch. The net result is that duration is preserved for all notes while pitch changes freely. An onset crossfade preserves the attack transient quality during the first few milliseconds of each note.

In **BOTH** mode the sample player decouples pitch from playback rate so that every triggered note plays at the same speed through the sample, regardless of which key on the keyboard triggered it. Pitch shifting is applied independently by the phase-vocoder time-stretch engine, while the source-read rate is driven solely by the BPM-stretch setting (or remains at one source sample per output sample when no BPM information is available). The visualiser cursor advances at the same rate for every voice, so a chord played in BOTH mode shows all of its voice cursors moving in lockstep across the waveform display rather than separating as they would under classical vinyl-style pitch shifting.

The SHIFT knob adds additional pitch shifting in semitones on top of the MIDI note offset, allowing creative transposition effects while maintaining constant duration.

2.7.2.8 PLAYBACK MODES IN INSTRUMENT MODE

Four of the six playback modes (OneShot, Loop Forward, Loop Backward and Pendulum) respond to MIDI notes and use the polyphonic voice system. Each MIDI note triggers a new voice that plays through its amplitude and filter envelopes independently.

Two modes bypass the voice system entirely and work the same as in the effect version:

Position mode -- The POSITION knob directly controls where in the sample audio is read from. MIDI notes have no effect. This mode is useful for scrubbing and modulation-driven playback.

Sync mode -- Playback is locked to the DAW transport and restarts at configurable beat intervals. MIDI notes have no effect.

In Position and Sync modes the amplitude envelope, filter, velocity curve and voice count controls are inactive since there are no MIDI-triggered voices to process.

2.7.3 MIDI MODULATION SOURCES

The instrument version adds four MIDI modulation sources to the modulation system alongside the existing LFOs, envelope followers and macros. These are available as modulation sources in the right-click modulation popup on any knob.

Pitch Bend - Responds to the MIDI pitch wheel. Outputs a bipolar signal from -1 to +1 (14-bit resolution).

Mod Wheel - Responds to MIDI CC #1 (modulation wheel). Outputs a bipolar signal mapped from the 0-127 MIDI range. A preset always loads as if the wheel were at zero, whatever position your wheel is physically in, so a preset never arrives with its mod wheel modulation already applied. It comes in as soon as you move the wheel or a CC #1 message arrives. This applies when you load a preset from the browser, with the preset arrows, with INIT and when a project or a DAW preset opens.

Aftertouch - Responds to MIDI channel pressure (aftertouch). Outputs a bipolar signal mapped from 0-127.

CC74 Slide - Responds to MIDI CC #74 which is the standard controller for MPE slide gestures on instruments like the Roli Seaboard and Sensel Morph. Outputs a bipolar signal mapped from 0-127.

Expression - Responds to MIDI CC #11, the standard expression pedal controller. Bipolar 0-127 mapping; routes that want unipolar pedal behaviour can toggle the route's BIP/UNI flag to remap the value into the 0..1 range.

Velocity - Sample-and-hold of the last received note-on velocity, mapped to the same bipolar -1..+1 range and centred on MIDI velocity 64 (the standard mezzo reference). A natural-intensity note produces zero offset; softer notes pull the modulated parameter below its base value and harder notes push above it. With the route in BIP mode this gives symmetrical "soft sounds darker, loud sounds brighter" behaviour around a neutral centre; with the route in UNI mode the value is remapped so the softest playable velocity contributes almost nothing and the loudest contributes a full modulation depth, which is the conventional "louder equals more open filter" pattern. The value is sticky between notes, so a single soft chord followed by silence keeps the parameter at the soft-chord level until the next note arrives.

When a Velocity route is set to BIP the source value behaves symmetrically around MIDI velocity 64, so soft notes pull the modulated parameter below its base value and hard notes push above it. When the same route is set to UNI the value is remapped into the 0..1 range, so the very softest playable velocity contributes almost no modulation and the loudest contributes a full depth - the conventional "louder equals more modulation" behaviour. Both modes draw from the same sample-and-hold value, so a single route can be reconfigured between the two without re-recording any automation.

All seven MIDI sources are smoothed with a 5 ms one-pole filter to prevent stepping artefacts from the 7-bit MIDI resolution. They appear in the modulation popup with a purple colour to distinguish them from LFOs (red), envelope followers (teal) and macros (amber).

2.7.4 REAL-TIME PARAMETER CHANGES

All instrument parameters (ADSR times, sustain levels, filter cutoff, resonance, key tracking and velocity curve) respond in real time to knob changes even while notes are sounding. If you turn up the amplitude sustain level while holding a chord the volume of the held notes increases immediately. If you shorten the filter decay while a note is in its decay stage the filter response accelerates to match the new setting. This makes the instrument version highly responsive to live performance adjustments and DAW automation.

2.7.5 PRESETS

Presets saved from the instrument version are tagged with `mode="instrument"` in the preset file. The preset browser automatically filters to show only instrument presets when running in instrument mode and only effect presets when running in effect mode. This prevents accidentally loading an instrument preset into the effect version (which would have no audio source) or vice versa.

3 PRESETS

3.1 UNDERSTANDING PRESETS

A preset in NeuralDream captures the complete state of the plugin - every layer, every engine, every parameter value, every modulation routing and every loaded sample. When you save a preset you are saving a snapshot of your entire processing configuration. When you load one you are restoring that exact configuration instantly.

Presets use the .ndpreset file format which is an XML document containing all parameter values, the layer topology (which engines are loaded in which order on which layers), modulation routing assignments and depths, preset metadata (name, category, notes) and any loaded audio.

One particularly powerful feature is that NeuralDream embeds loaded samples directly inside the preset file using lossless FLAC compression encoded as base64 text. This means your presets are fully portable - if you share a preset file with another user they will hear exactly what you hear even if they do not have the same sample files on their system. Impulse responses loaded into the IR Convolver engine are embedded in the same way.

3.2 PRESET CATEGORIES

NeuralDream organises presets into six categories:

Ambient - Evolving pads, atmospheric textures, drones and soundscapes. These presets typically use long reverb tails, granular processing and slow modulation to create immersive sonic environments.

Melodic - Presets designed for processing melodic content like vocals, synths and lead instruments. These tend to emphasise pitch-aware processing, harmonically tuned delays and musical filtering.

Rhythmic - Tempo-synced effects, gated patterns, rhythmic delays and pulsing modulations. These presets work best when NeuralDream's LFOs are locked to your DAW's transport.

Textural - Presets focused on transforming the timbral character of sound. Spectral processing, convolution, ring modulation and neural engines feature heavily in this category.

FX - One-shot effects, risers, impacts, glitch processing and dramatic transformations. These presets are designed for sound design and post-production work.

Utility - Practical processing tools like EQ curves, noise gates, stereo widening, simple delays and amp simulation setups. Starting points for building your own chains.

3.2.1 TAGS

A category answers one question and a preset can only have one. Tags answer the rest and a preset can carry as many as apply. Where the category says a preset is Ambient, its tags can say it is a pad, that it evolves, that it is warm and that it was made for scoring picture.

There are forty-five tags in four groups and one of those groups differs between the two plugins. SOUND says what kind of sound it is and in the instrument it offers PAD, PLUCK, LEAD, BASS, KEYS, STRINGS, BRASS, BELL, DRONE, TEXTURE, PERCUSSION, SOUNDFX, SUB, STAB, Mallet, VOX and RISER. In the effect plugin those words would mean nothing, so SOUND offers DISTORT, REVERB, COSMIC, ECHO, TWISTED, SMOOTH, PHASED, NEURAL, DEEP, STEREO-WIDE, FILTERED, COMPLEX, ETHEREAL, SYNC, VOCAL, SOLO and DRUMS instead.

The other three groups are the same in both plugins. BEHAVIOUR says what it does over time: MOTION, EVOLVING, STATIC, SEQUENCED, ARPEGGIATED, GENERATIVE, RHYTHMIC, GATED, GLIDE, SWEEP and WOBBLE. CHARACTER says how it sounds: WARM, BRIGHT, DARK, CLEAN, GRITTY, NOISY, GLASSY, ORGANIC, METALLIC, ANALOG and AGGRESSIVE. USE says what it is for: CINEMATIC, AMBIENT, TECHNO, LOFI, EXPERIMENTAL and DESIGN.

Because the SOUND words differ, a SOUND tag saved in one plugin is not part of the other's vocabulary. If you load an instrument preset in the effect or the other way round, its SOUND tag is not shown or searchable there, while its BEHAVIOUR, CHARACTER and USE tags still are.

The ENGINES column of the preset browser also has a GENERATORS group at its foot, holding Particles. Any preset with a layer whose source is Particles lights it, so you can find every particle patch in the library with one click, exactly as you would filter by any engine.

The list is fixed on purpose. A library where one preset is tagged AMBIENT and another Ambient is worse than no tags at all, so tags are chosen by clicking rather than typed and anything not in the list is never stored. In the browser, clicking tags narrows the list to presets that carry all of them. Presets saved before tags existed simply have none and gain them the next time you save them.

3.3 SAVING PRESETS

To save your current configuration as a preset click the SAVE button in the preset strip or the SAVE button in the top right of the preset browser. The browser opens in save mode: the same panel you browse in, doing a different job.

The CATEGORY column becomes the chooser and you click one of the six. The TAGS column becomes this preset's tags and you click as many as apply; nothing is typed, so there is no vocabulary to remember and no way to invent a misspelled tag. On the right are three fields: PRESET NAME, AUTHOR and DESCRIPTION. The description is where you say what the preset is and how to play it and it is what the browser shows in the preview strip and in the search. The author is offered pre-filled with whatever you entered last time, so it is typed once and then simply confirmed. Along the foot the panel reads back the category and the tags you have chosen, so they can be checked before committing. SAVE PRESET writes it; CANCEL returns to browsing.

Your preset is saved to the user presets directory and appears immediately in the browser. Each preset gets its own .ndpreset file named after the preset name you chose. If the preset you are trying to save has the same name as a preset already on your filesystem then NeuralDream will prompt you to confirm that you wish to overwrite it.

3.4 LOADING PRESETS

There are three ways to load a preset:

- ★ Use the left and right Chevrons either side of the centred and loaded Preset Title in the preset strip to step through presets one at a time. This is the quickest way to audition presets in sequence.
- ★ Click the preset name to open the preset browser, where you can search, sort, filter by category, engine or tag, read each preset's description and load with a click, a double click or the arrow keys.
- ★ Drag a .ndpreset file from Finder directly onto the plugin window.

Loading a preset replaces your current state entirely. If you want to keep your current work use SAVE first or rely on the undo system to step back after auditioning.

3.5 FILE LOCATIONS

On macOS factory presets are installed to /Library/Application Support/DreamDeep/NeuralDream/ by the installer. These are read-only - you cannot overwrite them from within the plugin. User presets are saved to:

~/Library/Application Support/DreamDeep/NeuralDream

On Windows factory presets are installed to C:\ProgramData\DreamDeep\NeuralDream\ by the installer. These are read-only - you cannot overwrite them from within the plugin. User presets are saved to:

C:\Users\<your name>\AppData\Roaming\DreamDeep\NeuralDream

If a user preset has the same filename as a factory preset then the user version takes priority. This cascading system lets you create personalised variations of factory sounds. The original factory version remains untouched on disk and can be restored simply by deleting your user version.

3.6 SUB-PRESETS

In addition to full plugin presets NeuralDream supports sub-presets - saved configurations for individual engines. Each engine panel has its own sub-preset controls in the header bar allowing you to save and recall just that engine's settings independently.

Sub-presets are useful when you have dialled in a particular sound on one engine and want to reuse it in a different context. For example you might save a carefully tuned Parametric EQ curve as a sub-preset and then recall it across multiple full presets without having to recreate it each time.

Sub-presets are stored as standard .ndpreset files and are scanned from the same preset directories as full presets. The sub-preset browser filters the list to show only presets that contain the relevant engine type.

3.7 EXPORTING PRESETS

The preset browser includes an Export button at the bottom of the category panel. Clicking it opens a folder chooser dialog (defaulting to your Desktop) where you select a destination folder. NeuralDream then copies all of your user presets to that folder and displays a confirmation showing how many presets were exported.

This is useful for backing up your presets, sharing them with other users or transferring them to another machine. The exported files are standard .ndpreset files with all parameter values,

modulation routings and embedded audio intact - the recipient can simply drop them into their own NeuralDream presets directory and they will appear in the browser.

4 THE LAYER SYSTEM

4.1 WHAT ARE LAYERS?

Layers are the fundamental building blocks of NeuralDream. Each layer is an independent audio processing path with its own audio source, its own chain of transformer engines and its own output controls. Think of each layer as a separate effects rack that happens to live inside the same plugin instance.

The power of layers comes from combining them. Run the same input through two different processing chains and blend the results. Feed one layer's output into another to create serial processing stages. Use a sample on one layer as a modulator for a vocoder on another. The architecture is deliberately open-ended - there is no single correct way to use layers and some of the most interesting sounds come from unexpected routing combinations.

NeuralDream supports up to eight simultaneous layers. Their outputs are summed together to produce the plugin's final output with each layer contributing according to its own mix and gain settings.

4.2 AUDIO SOURCES

Each layer draws audio from one of several source types selected via the source dropdown on the top right-hand side of the layer strip:

Plugin Input - The audio arriving from your DAW. This is the default source for the first layer and the most common starting point. Whatever audio track NeuralDream is inserted on provides this input.

Sidechain - the audio arriving on NeuralDream's sidechain input from your DAW. You route a different track into it in the host, then use it as a layer's source or as the modifier for a dual-source engine, which is how you make one sound duck, gate or shape another.

Sample - A loaded audio file played back by NeuralDream's built-in sample player. When you select Sample as a source the sample player controls appear allowing you to load a file, set loop points, adjust playback speed, enable time stretching and more. See Chapter 5 for full details on the sample player.

Noise - A built-in coloured noise generator with two controls. The COLOUR knob sweeps continuously from Brown through Pink, White, Blue and Violet noise. Each noise colour has different spectral characteristics - brown noise is warm and bassy, pink noise has equal energy per octave, white noise is flat across all frequencies, blue noise emphasises high frequencies and violet noise is the brightest. The LEVEL knob sets the noise output level. Noise is useful as a source for filtering, gating and spectral processing or as a textural layer blended underneath other material.

Particles - a sound generation engine of its own and the newest source in NeuralDream. It does not play back a sample and it is not the synthesiser. It emits a stream of very short, independent tones, each one born with its own pitch, wavetable, length, loudness and stereo position, each carrying its own envelopes and low-frequency oscillators and each dying on its own schedule. It is available on the effect build as well as the instrument, because it can run without any MIDI at all. It arrives with TRACK switched on, which means it stays silent until you play a note, so adding it to a layer never takes you by surprise. Switch TRACK off in the PITCH section and it free-runs, emitting continuously whether or not anything is played. Selecting it opens a full-width control panel, the same size as the synthesiser's. See Chapter 12 for the whole engine.

Layer N - the output of another layer or a single oscillator within another layer's synthesiser. Alongside the plain "Layer N" entries the source menu offers per-oscillator taps - for example "Layer 1 → OSC 2" - which draw just that oscillator's signal, letting one synthesiser voice be processed in isolation by a whole chain of its own. Only layers with a lower index can be selected as sources, preventing feedback loops.

An important detail about both kinds of tap: the audio is taken before the mixer and output stages of the source layer. A whole-layer tap is drawn before that layer's output LEVEL and pan. A per-oscillator tap is the oscillator as you would hear it with its fader up: it has passed through the filter it is routed to, both filters in series in the SERIAL topology and through the note's VCA envelope but not through the mixer channel's level or pan. This is deliberate - the mixer and fader govern what a layer contributes to its own output, never what other layers may borrow from it. You can therefore turn an oscillator fully down in its own mixer, yet still feed it, filtered and enveloped, at full strength into another layer for independent processing.

4.3 THE TRANSFORMER CHAIN

Each layer can host one or more transformer engines in series. Audio flows through the chain from left to right - the output of the first engine feeds the input of the second and so on. The final engine's output becomes the layer's processed signal.

To add an engine click the + button at the right end of the chain. A categorised menu appears showing all available engines grouped by type - Single Source DSP (blue), Single Source ML (purple), Dual Source DSP (teal) and Dual Source ML (amber). Select an engine and it appears in the chain with its accent colour.

Engines can be reordered by dragging the handle icon in the top-left corner of each engine block. The order matters - a reverb followed by a delay sounds very different from a delay followed by a reverb.

To remove an engine right-click its block and select Delete from the context menu. The remaining engines close the gap automatically.

There is no fixed limit on the number of engines per layer but each engine adds CPU load so complex chains on multiple layers will demand more processing power.

Some engines need a moment to do their work, so they hand their sound back slightly later than they received it. These include the Exciter, Aperture, Resonance Field, the spectral engines, RAVE and a time-stretched sample. NeuralDream never lets a delayed signal meet an undelayed one. Wherever two signals come together, the earlier one is held back until it lines up exactly with the later one. That happens in four places: in each engine's MIX, in each engine's bypass, at the second input of a dual source engine and where the layers are added together, where every layer is lined up with the slowest one. So whatever you set and whichever engines you load, nothing is mixed out of step. The one deliberate exception is a dual source engine whose second input is another layer's output: as section 4.4.1 describes, that input has always arrived one audio block later, by design.

NeuralDream then tells your DAW how late its output is: the total for the slowest layer, which includes its source and every engine in its chain, because engines in a layer run one after another and their delays add up. Your DAW uses this figure to keep NeuralDream in time with your other tracks and NeuralDream updates it straight away whenever you add, remove, replace or move an engine and whenever you load a preset or a project. When no engine that needs extra time is loaded, nothing is held back and there is no delay at all. 4.4 Dual Source Engines

Six of NeuralDream's engines are dual source - they require two audio inputs rather than one. These are Spectral Morph, Spectral Gate, Vocoder, Live Convolve, Ring Modulator and Spectral Style Transfer.

When you add a dual source engine to a chain it processes the layer's audio as its primary (carrier) signal. The secondary (modifier) signal is configured separately via a dropdown on the engine's panel. The modifier can be drawn from the plugin input, a dedicated sample player or

the output of another layer - giving each dual source engine its own independent second audio feed.

Each dual source engine's modifier has its own sample player instance with full controls for loading files, setting loop points and adjusting playback. This means you can have a vocal sample modulating a vocoder on one engine while a completely different sample drives a spectral morph on another engine in the same chain.

4.4.1 MODIFIER ROUTING OPTIONS

The Modifier dropdown menu on each dual-source engine gives options Plugin Input / Sidechain / Sample / Layer Output 1, 2, 3, etc). You can't pick a layer whose modifier feed would loop back.

The layer-output modifier feed is taken pre-fader so muting/soloing the source layer for monitoring purposes does not silence the modulation feeding the dual-source engine.

A use-case example (e.g. "set Vocoder modifier to Layer 1 output so a layer of synth pads becomes the carrier for the vocoder while another layer's vocal sample becomes the modulator").

4.5 LAYER OUTPUT CONTROLS

Each layer strip has output controls on its right side:

MIX - The dry/wet balance for the entire layer. At 0% you hear only the unprocessed source audio. At 100% you hear only the output of the transformer chain. Values in between blend the two. This is applied after the engine chain so it affects all engines simultaneously.

GAIN - The output level for the layer in dB. This controls how loud the layer's contribution is in the final mix. Level changes are smoothed with a one-pole exponential filter (~5ms time constant) to prevent clicks and zipper noise when adjusting during playback.

Solo (S) - Solos this layer, muting all others. When multiple layers have Solo engaged you hear all of the soloed layers together. Solo is useful for isolating a single processing path to fine-tune it without distraction.

Mute (M) - Silences the layer's contribution to the final output. The layer continues to process audio internally so unmute restores the sound instantly without any re-initialisation delay.

Bypass (B) - Bypasses all engines in the layer's chain, passing the source audio straight through to the output. Unlike Mute the audio still reaches the output - you hear the dry source signal. This is useful for A/B comparison between processed and unprocessed sound.

4.6 ADDING AND REMOVING LAYERS

Click the + button below the layer area to add a new layer. It appears at the bottom of the stack with Plugin Input as its default source and an empty transformer chain ready for you to add engines.

To remove a layer right-click its header strip and select Delete. A confirmation is not required - but undo is available if you delete a layer by accident. The first layer cannot be deleted as NeuralDream always requires at least one active layer.

Layers can be reordered by dragging the handle on the left side of each layer strip. The order affects inter-layer routing - remember that a layer can only draw audio from layers above it in the stack.

5 THE SAMPLE PLAYER

5.1 OVERVIEW

NeuralDream includes a fully featured sample player that can be used as an audio source for any layer. Rather than simply looping a file the sample player offers six distinct playback modes, adjustable loop points with zero-crossing snap, independent time stretching and pitch shifting, interactive waveform navigation and automatic BPM and root note detection. Samples are embedded directly into presets using lossless FLAC compression so your presets are fully portable.

The sample player supports WAV, AIFF, FLAC, OGG, MP3 and AAC files up to 10 minutes in length. Files are resampled to match your DAW's session sample rate on load and peak-normalised to 0.99 to ensure consistent levels across different source files.

5.2 LOADING SAMPLES

To use a sample as a layer source select Sample from the layer's source dropdown. The sample player controls appear showing the waveform display and control panel.

To load a sample click the LOAD button or drag an audio file from Finder directly onto the waveform area. The file is loaded, resampled to the session sample rate if necessary and displayed as a waveform. Loading happens off the audio thread so there is no interruption to playback.

When a sample is loaded NeuralDream automatically attempts to detect the original BPM and root note from the filename. Common naming conventions like "120bpm" or "C#min" are recognised. If BPM cannot be determined from the filename NeuralDream estimates it by assuming the file represents 1, 2, 4 or 8 bars of 4/4 time. The detected values are displayed in the info area and used by the time stretch engine when enabled.

5.3 THE WAVEFORM DISPLAY

The waveform display shows your loaded sample as a mirrored amplitude view with markers for the start position, end position, loop start, loop end and playhead. Each marker type has a distinct colour - green for start, red for end, blue for loop boundaries and white for the playhead.

The waveform supports interactive zoom and pan:

Mouse wheel zooms in and out from 1x to 64x magnification. Zooming centres on the mouse cursor position so you can zoom directly into the region you are interested in.

Shift + mouse wheel pans horizontally through the waveform when zoomed in.

Apple Mouse side-swipe also pans horizontally for trackpad users.

When zoomed in the waveform automatically scrolls to follow the playhead during playback so you never lose sight of the current playback position.

The grid overlay adapts to your zoom level showing quarter notes at low zoom, then progressively finer divisions (8ths, 16ths, 32nds, 64ths) as you zoom in. When time stretch is active the grid uses the original BPM of the sample for accurate beat-aligned editing. Marker snapping follows the visible grid resolution with a minimum of 16th note resolution.

Start and end markers can be dragged to trim the playback region. NeuralDream automatically snaps markers to zero-crossing points within a window of 512 samples to avoid clicks at the boundaries.

5.4 PLAYBACK MODES

The MODE button on the control panel cycles through six playback modes, each with its own colour indicator:

OneShot (grey) - Plays the sample once from the start position to the end position then stops. Useful for one-shot sound effects, hits and stabs. Playback begins when the audio level at the trigger input crosses the 50% threshold.

Loop Forward (green) - Plays from the sample start position to the loop end then jumps back to the loop start and repeats continuously. This is the most common mode for sustaining pads, textures and rhythmic loops. A crossfade control smooths the loop transition to eliminate clicks at the loop point.

Loop Backward (orange) - Plays from the sample end position backward to the loop start then jumps back to the loop end and repeats. This produces a reversed loop effect while keeping the loop boundaries the same as Loop Forward.

Pendulum (violet) - Plays from the sample start position to the loop end position and then plays forward to the loop end, reverses to the loop start, plays forward again and so on. This creates smooth palindromic loops that avoid any discontinuity at the loop points.

Position (teal) - The playback position is controlled directly by the POSITION knob rather than advancing automatically. This turns the sample into a scrubbing instrument - automate or modulate the position knob to scan through the audio at any speed and in any direction. This mode is particularly effective when driven by an LFO for rhythmic granular-like effects or by an envelope follower for dynamic scrubbing.

Sync (cyan) - Playback is locked to your DAW's transport. The sample restarts at configurable intervals synchronised to the beat grid - 4 bars, 2 bars, 1 bar, 1/2, 1/4, 1/8, 1/16 or 1/32. This ensures the sample stays in time with your project regardless of tempo changes.

5.5 LOOP CONTROLS

The control panel provides dedicated knobs for fine-tuning loop behaviour:

START - Sets the playback start position within the sample.

END - Sets the playback end position. Audio outside the start/end range is never played.

LP START - Sets the loop start point (used by Loop Forward, Loop Backward and Pendulum modes). This can be placed anywhere between START and END.

LP END - Sets the loop end point. Together with LP START this defines the looping region.

XFADE - Controls the crossfade length at the loop point. Higher values create a smoother transition at the cost of a slightly less precise loop boundary. This is essential for eliminating clicks in loops that do not have natural zero-crossing points at their boundaries.

SPEED - Adjusts the playback rate of the sample. At 100% the sample plays at its original speed. Lower values slow it down and higher values speed it up. This changes both the speed and pitch of the audio (unless time stretch is enabled).

5.6 TIME STRETCHING

The STRETCH button enables real-time time stretching which allows the sample to match your DAW's tempo without changing its pitch. When stretch is active the sample's playback speed is adjusted automatically based on the ratio between its original BPM and the current DAW tempo.

NeuralDream's time stretch engine uses a phase vocoder with identity phase locking following Laroche and Dolson (1999) to maintain phase coherence across frequency bins. Transient detection based on Roebel (2003) identifies attacks and drum hits in the audio and resets the phase at those points to preserve the crispness of percussive material.

The stretch engine uses an FFT size of 4096 with 4x overlap for high-quality results. This introduces approximately 93ms of latency which is compensated automatically by the plugin.

The info area on the control panel displays three values:

ORIG - The detected or manually set original BPM of the sample.

DAW - The current tempo reported by your DAW.

ROOT - The detected root note of the sample if one was found in the filename.

Time stretching is supported in OneShot, Loop Forward and Sync playback modes.

5.7 PITCH SHIFTING

The SHIFT knob provides independent pitch shifting in semitone steps from -24 to +24 (two octaves down to two octaves up). Pitch shifting works independently of time stretching - you can change the pitch without affecting playback speed and vice versa or use both simultaneously.

The pitch shift is implemented using spectral bin resampling within the same STFT framework as the time stretch engine so both features share the same high-quality phase vocoder processing.

5.8 WAVEFORM OVERLAY

The sample control panel features a subtle waveform overlay drawn behind all the controls. This shows the full shape of the loaded sample at a glance using the sample player's accent colour at 70% opacity. The overlay updates automatically when a new sample is loaded or when the panel is resized.

6 AUDIO TRANSFORMERS

6.1 ENGINE OVERVIEW

NeuralDream's 40 (or so) transformer engines are the heart of the plugin. Each engine is a self-contained audio processor with its own controls, its own visualiser and its own accent colour. Engines are added to a layer's transformer chain and process audio in series from left to right.

Engines fall into four categories visible in the add-engine menu (the triple-dot button that sits at the end of the existing engines in each layer):

Single Source DSP (blue) - Conventional signal processing engines that take one audio input. This is the largest category covering filters, delays, reverbs, modulation effects, dynamics and more.

Single Source ML (purple) - Machine learning engines that take one audio input and process it through a neural network.

Dual Source DSP (teal) - Engines that require two audio inputs - a carrier and a modifier. These include spectral processors, vocoders and ring modulation.

Dual Source ML (amber) - Machine learning engines that require two audio inputs. RAVE operates in this category when using sidechain audio as the modifier.

The following sections describe each engine in detail including its purpose, its parameters and any special features.



Your DAW sees the same 1024 slots as automation parameters. A slot in use is named for what it controls, so a delay added to Layer 2 appears as L2 Delay Feedback, L2 Delay Time and so on and the names update as engines are added, removed or swapped. A slot not in use is named (unused) followed by its number, so the live controls stand out from the filler when you scroll the list. Automation recorded against an engine that is later removed is simply ignored.

One thing in NeuralDream is deliberately not in this chapter. PARTICLES is not an audio transformer: it does not process audio arriving from somewhere else, it generates sound of its own, so it is a layer SOURCE rather than an engine in the chain. It is described in Chapter 12.

6.2 LP/HP FILTER

A dual-mode resonant filter with selectable roll-off slope. The engine operates as either a low-pass or a high-pass filter and can be configured from a gentle two-pole curve to an aggressive eight-pole cut. Implemented as a cascade of topology-preserving transform (TPT) state-variable filter stages for numerical stability across the full frequency range even under heavy modulation.

Unlike the plateau shelves of the Parametric EQ this filter continues to roll off past the corner frequency - the deeper the slope the more surgical the cut. Useful for carving space out of a mix, removing unwanted low-end rumble or high-frequency hiss or creating classic filter-sweep effects under modulation.

MODE (LP / HP) - Selects low-pass or high-pass operation. In LP mode the filter attenuates frequencies above the cutoff. In HP mode it attenuates frequencies below. The mode switch sits above the CUTOFF knob in the panel.

SLOPE (12 / 24 / 36 / 48 dB) - Selects the filter steepness in decibels per octave. Each step adds a cascaded two-pole stage. 12 dB/oct is a single stage with a gentle natural curve, 24 dB/oct is the default and a musical all-rounder, 36 dB/oct is aggressive and surgical and 48 dB/oct is the steepest setting for brick-wall-style cuts. The slope selector sits above the RESONANCE knob in the panel.

CUTOFF - Sets the filter cutoff frequency from 20 Hz to 20 kHz. Fully modulatable - assign an LFO or envelope follower for classic filter sweeps and auto-wah effects.

RESONANCE - Adds emphasis at the cutoff frequency from a smooth roll-off at 0% up to self-oscillation at higher settings. To keep the resonant character consistent across slope settings resonance is applied only to the final stage of the cascade - the earlier stages run flat. This is the standard technique for resonant cascaded filters and prevents excessive peak stacking when multiple resonant stages are in series.

The visualiser shows the filter's frequency response curve across the full audible range. The curve updates live with mode, slope, cutoff and resonance changes so you can see exactly how the filter is shaping your signal. A vertical marker indicates the current cutoff frequency.

6.3 VOLTAGE-CONTROLLED FILTER (VCF)

A more characterful filter inspired by analog synthesiser VCF circuits. The VCF uses a topology-preserving transform state variable filter with per-sample smoothing in the log-frequency domain following the approach described by Vadim Zavalishin in "The Art of VA Filter Design." The log-domain smoothing ensures that sweeping the cutoff across wide ranges - especially under LFO modulation - produces a smooth musical sweep without clicks or zipper noise.

CUTOFF - Filter cutoff frequency from 20 Hz to 20 kHz.

RESONANCE - Feedback emphasis at the cutoff frequency.

MODE - Selects between five filter types: LP (low pass), HP (high pass), BP (band pass), NOTCH (band reject) and COMB (comb filter). Each mode has a fundamentally different character - low pass is warm and subtractive, band pass isolates a frequency region, notch removes a narrow band and comb creates a series of harmonically spaced peaks and nulls.

SLOPE - Selects the filter steepness: 6dB, 12dB or 24dB per octave. The 24dB setting uses a circuit emulation of a famous Ladder filter topology for a thick resonant character. Slope buttons are greyed out in COMB mode as the comb filter uses a fixed topology.

The VCF is particularly well suited to modulation. Assigning an LFO or envelope follower to the cutoff produces the classic auto-wah and filter sweep effects heard in electronic music. The NeuralDream Instrument provides many filter Models and (sub) Types. See section 11.3.5.1 for more information.

6.4 PARAMETRIC EQ

A three-band parametric equaliser using JUCE IIR biquad filters based on the RBJ Audio EQ Cookbook. Each band can boost or cut by up to 24 dB with adjustable centre frequency and Q (bandwidth).

LO FREQ - Centre frequency of the low band from 20 Hz to 2 kHz. This band operates as a low shelf.

LO Q - Bandwidth of the low band.

LO GAIN - Boost or cut for the low band from -24 to +24 dB.

MID FREQ - Centre frequency of the mid band from 100 Hz to 10 kHz. This band operates as a peaking bell.

MID Q - Bandwidth of the mid band. Lower values give a broad gentle curve while higher values give a narrow surgical cut or boost.

MID GAIN - Boost or cut for the mid band from -24 to +24 dB.

HI FREQ - Centre frequency of the high band from 1 kHz to 20 kHz. This band operates as a high shelf.

HI Q - Bandwidth of the high band.

HI GAIN - Boost or cut for the high band from -24 to +24 dB.

Filter coefficients are interpolated per-sample with a one-pole smoother (~2ms time constant) to eliminate clicks when parameters change rapidly under modulation. The visualiser shows:

- ★ A 512-point logarithmic frequency response plot with three coloured curves - red for the low band, gold for the mid band and blue for the high band - each with a gradient fill
- ★ The FFT spectrum analyser overlay (added 1.1.24) drawing a translucent blue-cool filled shape behind the EQ curve showing the input signal spectrum
- ★ The PRE/POST toggle pill in the bottom-right corner for switching the spectrum tap between pre-EQ and post-EQ.

6.5 NOISE GATE

A feedforward noise gate following the topology described by Giannoulis, Massberg and Reiss (2012) in "Digital Dynamic Range Compressor Design - A Tutorial and Analysis" published in JAES Vol. 60 No. 6. The gate attenuates audio that falls below a threshold, useful for removing noise, bleed and unwanted low-level sound between phrases.

THRESHOLD - The level in dB below which the gate closes. Audio above this level passes through unchanged.

ATTACK - How quickly the gate opens when the signal exceeds the threshold, from 0.1 to 50 ms. Short attack times let transients through cleanly while longer times produce a softer fade-in.

HOLD - How long the gate stays open after the signal drops below the threshold, from 0 to 500 ms. This prevents the gate from chattering on and off during sustained sounds that hover near the threshold.

RELEASE - How quickly the gate closes after the hold time expires, from 5 to 2000 ms. Short release times cut the tail abruptly while longer times produce a natural fade-out.

RANGE - The depth of attenuation when the gate is closed, from 0 to 100%. At 100% the gate silences the audio completely. At lower values some signal bleeds through even when the gate is closed which can sound more natural.

SOURCE - Detector input selector. Chooses which audio signal drives the gate's level detector while the gate itself always processes the layer's own audio. Options are Input (the plugin's main input), Layer 1 through Layer 8 (the output of any other layer in the patch) and Sidechain (the DAW sidechain bus). Use this to make one layer's audio open and close in time with another layer's dynamics - for example, gating a pad layer from a drum layer so the pad only sounds during drum hits or gating a reverb send against a vocal so the verb breathes with the singer's phrasing.

The gate uses 3 dB of hysteresis between the open and close thresholds to prevent chattering at the boundary. Smoothing coefficients are derived directly from the Giannoulis paper (Equations 7 and 8). The visualiser shows a scrolling level history with a filled region indicating gate state (open or closed), a dashed threshold line with dB label, an OPEN/CLOSED text indicator and a gain reduction meter on the right edge.

6.6 DELAY

A multi-tap tempo-syncable delay with feedback, stereo panorama and saturation.

TIME - The delay time in milliseconds (free mode) or beat divisions (synced mode).

FEEDBACK - Controls how much of the delayed signal is fed back into the delay input. Higher values produce more repeats.

SMEAR - Cross-feed between delay taps creating a diffuse, reverb-like decay quality rather than distinct rhythmic echoes.

PANORAMA - Chaos-modulated stereo width. Adds random stereo movement to the delay taps creating a wider, more spatially interesting echo field.

SUP/COM - Sidechain dynamics control. Negative values duck the delay output when the dry signal is loud (sidechain ducking). Positive values compress the delay output against the dry signal.

SPARKLE - Presence and air boost applied to the delay output giving echoes a bright, airy quality that helps them sit above the dry sound.

TAPS - Number of delay taps from 1 to 4. More taps create denser rhythmic patterns. If you are not in SYNC mode then these taps follow patterns of prime numbers, ensuring that they never quite repeat - Ambient Delight!

SYNC - Locks the delay time to your DAW's tempo.

DIVISION - Beat division for tempo sync when SYNC is enabled.

6.7 STEREO ECHO

A true stereo delay engine with two independent Hermite-interpolated delay lines, one for each channel. This engine offers far more control than the basic Delay and can produce everything from subtle stereo widening to complex ping-pong patterns.

TIME L / TIME R - Independent delay times for the left and right channels up to 2 seconds each. These can be set freely or synced to tempo with independent beat divisions.

FEEDBACK L / FEEDBACK R - Independent feedback amounts for each channel.

XFEED - Cross-feed between the two delay lines. At 0% the delays are fully independent (true stereo). At 100% the output of each delay feeds into the opposite channel creating a ping-pong effect that bounces between left and right.

LO CUT - A four-pole 24 dB/oct high-pass filter applied to all echo output including feedback. Removes low-frequency buildup in the echoes keeping the delays clean and defined.

HI CUT - A four-pole 24 dB/oct low-pass filter applied to all echo output including feedback. Darkens successive echoes for a natural tape-echo quality.

DRIVE - Asymmetric soft clipping applied to the echo output (not in the feedback path) with pre/de-emphasis and output compensation. Adds warmth and saturation to the delayed signal.

SYNC L / SYNC R - Independent tempo-sync toggles for the left and right delay lines. Either channel can be free-running in milliseconds while the other is locked to the DAW tempo, allowing for hybrid free-versus-synced echo patterns that pure stereo delays cannot produce.

BEAT L / BEAT R - Independent beat-division selectors for the left and right channels when their respective Sync toggle is on. Each channel can lock to a different note value, so for example the left tap can hit on every dotted eighth while the right tap hits on every quarter - a classic stereo-rhythmic-displacement effect.

The visualiser shows a tap timeline with hot pink markers for the left channel and electric cyan for the right channel.

6.8 FDN REVERB

An eight-line feedback delay network. Each line is a short delay, between about 15 and 35 milliseconds at the smallest SIZE and between about 240 and 560 milliseconds at the largest. The eight lengths are mutually prime, so their echoes never pile up on the same moments. The sound circulates through the lines and on every trip a mixing network passes it from each line to the others. At WEAVE 100 that network is the Hadamard matrix of Jot and Chaigne (1991), "Digital Delay Networks for Designing Artificial Reverberators", from the Proceedings of the AES Convention, which spreads every echo across all eight lines and produces dense, smooth reverberation.

The length of the tail is set as a time rather than as a feedback amount. Each line is given exactly the feedback that makes it die away at the chosen rate, so every part of the tail fades together and the reverb stays clear rather than muddy. The engine runs as two complete networks, one for the left input and one for the right, so it is stereo from input to output and the reverb returns at about the same level as the signal going in.

The controls are described below in the order they appear on the panel, reading across each row.

SMEAR - Controls how much the incoming sound is blurred before it enters the network and how far each line drifts slowly in length. Higher values smooth the first reflections into a soft wash and add gentle movement to the tail, while lower values keep attacks crisp. The blurring of the

input follows WEAVE, so as WEAVE comes down the blurring fades with it and the separate echoes stay clean.

PANORAMA - Controls the width and movement of the reverb's stereo field. At 0% each reverb line sits at a fixed position, giving a still, conventional stereo image. Turning PANORAMA up collapses the fixed positions inward and introduces chaos-driven motion so the lines drift unpredictably across the stereo field. Low settings add subtle life and high settings produce an obvious swirling motion that is especially effective on stereo sources.

SIZE - Controls the dimensions of the space. At 0% the lines are between about 15 and 35 milliseconds long, which is a tight room and at 100% they are between about 240 and 560 milliseconds, which is a vast one. SIZE changes only the spacing of the reflections, not how long the tail lasts, so use DECAY for that. When WEAVE is turned down, SIZE also sets the repeat time of the echoes.

WEAVE - Sets how tightly the eight lines are woven together and it turns the engine from a reverb into an echo. At 100% every line feeds every other line on each trip and each line keeps its own length, which gives a full, dense reverb. As WEAVE comes down, the lines are woven together less and their lengths glide towards one shared length, so the sound thins from a reverb into a cluster of repeats. At 0% the eight lines repeat together as one clean echo and SIZE sets its repeat time: about 115 milliseconds at SIZE 50 and about 390 milliseconds at SIZE 100. Settings in between are neither reverb nor echo and that is where the most unusual sounds are found. Turning WEAVE while a sound is still ringing bends its pitch briefly, in the same way that changing the time on a tape delay does.

SHIMMER - Adds two pitch-shifted voices to the tail, one an octave above and one an octave and a fifth above. The voices are taken from the reverb itself and fed back into it, so every pass climbs higher and the tail rises into a glittering upper register. The second voice enters a fraction later than the first, so the two overlap rather than starting together and the pair drift very slightly in pitch and level against each other, so they sound alive rather than mechanical. Each time the shimmer returns it is kept between 120 Hz and 7.5 kHz, so the rising stack fades away at the top instead of building forever. The strength of the shimmer follows DECAY and is limited so that the reverb can never run away at any setting of any control. At the very top of DECAY, where the tail holds, there is no room left for the shimmer and it falls silent. At 0% the shimmer is off.

WANDER - Adds a slow random drift to the pitch of the two shimmer voices, up to 18 cents either way at 100%, so the octaves are not surgically exact. The left and right sides drift independently of each other. WANDER only acts on the shimmer, so it has no effect when SHIMMER is at 0%.

SKEW - Makes the room ring for longer on one side than the other, from -100% to +100%. Turned clockwise, the left side's tail lengthens while the right side's shortens and each speaker leans further towards its own side of the reverb. Turned anticlockwise, the reverse happens. At DECAY 50 with SKEW fully clockwise, the left side takes about 4.5 seconds to fall by 30 dB while the right side takes about 0.3 seconds, so one side of the room rings on after the other has gone quiet. At full SKEW, a source panned hard to one side keeps almost all of its reverb on that side. At 0 both sides decay together.

DECAY - Sets how long the tail lasts, independently of SIZE. At 0% the tail dies away in about a third of a second, at 50% in about three seconds and at 90% in about fifteen seconds. At the very top of its travel the tail holds almost indefinitely. Reach for DECAY when a preset sounds right but the tail is wrong for the section.

LOW - A low shelf at 350 Hz inside the reverb, from -18 to +9 dB, so it shapes the tail a little more on every pass. Cutting it thins the low end of the tail and boosting it adds warmth and weight. A boost is balanced by a slightly shorter tail at the low end, which is what stops the reverb from ever running away.

HIGH - A high shelf at 2 kHz inside the reverb, from -18 to +12 dB, so it shapes the tail a little more on every pass. Cutting it gives the soft, darkening tail of a room with absorbent surfaces and boosting it gives a brighter, more present tail. A boost is balanced by a slightly shorter tail at the high end, which is what stops the reverb from ever running away.

REVERSE - Toggles reverse reverb mode, where the reverb tail builds up before the transient rather than decaying after it.

6.9 SUPERSPACE

A large-scale reverb/delay hybrid inspired by Valhalla SuperMassive. Five distinct algorithms produce everything from lush reverbs to cascading echoes and dense granular clouds.

MODE - Selects between five algorithms: GEM, HYD, SAG, AQU and PLE. Each has a different delay line topology and feedback structure. The visualiser changes colour palette for each mode.

DELAY - Base delay time in milliseconds. This can be tempo-synced using the SYNC toggle.

WARP - Delay time spread. Introduces nonlinear warping of the delay times creating pitch-shifting and detuning effects within the feedback network.

FEEDBACK - Regeneration amount. Controls the overall feedback. High values produce infinite sustain.

DENSITY - Inter-delay mixing. Controls how quickly the sound becomes dense and diffuse.

LOW - Low frequency damping. Sets the frequency below which the reverb tail is attenuated faster.

HIGH - High frequency damping. Sets the frequency above which the reverb tail is attenuated faster.

MOD - Modulation depth. Adds subtle pitch modulation to the delay lines reducing metallic ringing and adding a lush chorusing quality.

MOD RT - Modulation rate. Controls the speed of the internal modulation.

WIDTH - Stereo width from -100% to +100%.

SYNC - Locks the delay time to your DAW's tempo.

DIVISION - Beat division for tempo sync.

6.10 DARKSKY

An atmospheric diffuse reverb with shimmer inspired by the Strymon NightSky pedal. Designed for creating vast ambient spaces, drones and evolving textures rather than realistic room simulation.

DECAY - Reverb decay time from short ambience to endless sustain.

SIZE - The reverb space from intimate to infinite.

SHIMMER - Octave-up pitch shifting in the feedback path.

TONE - Tonal colour from dark to bright, shaping the frequency balance of the reverb tail.

TEXTURE - Diffusion texture. Controls the internal diffusion character of the reverb.

MOD - Modulation depth adding movement and chorus-like width to the tail.

LOW - Low shelf EQ from -18 to +9 dB.

HIGH - High shelf EQ from -18 to +9 dB.

MOD RT - Modulation rate controlling the speed of the internal modulation.

The visualiser shows a 3D perspective aurora with flowing curtains of colour - bright mint at the horizon rising through electric green and vivid gold to hot pink overhead. The aurora responds to the reverb's energy level and spectral content.

6.11 CHORUS

A stereo chorus based on the bucket-brigade device (BBD) delay line model with modulated delay times creating the characteristic ensemble thickening effect of analog chorus units. This engine goes beyond basic chorus with WIDE and HYPER modes for extreme stereo enhancement.

RATE - The speed of the LFO modulating the delay time.

DEPTH - How far the delay time is modulated. Higher values produce a more dramatic pitch wobble.

WIDTH - Stereo width of the voice panning.

STONE - BBD lowpass filter brightness. Shapes the tonal character of the chorus effect simulating the natural high-frequency rolloff of analog bucket-brigade circuits.

WIDE - A two-part control. The WIDE toggle activates stereo decorrelation mode and the WIDE depth knob controls how much decorrelation is applied. WIDE mode adds phase-offset copies of the chorus voices for a dramatically wider stereo image.

HYPER - A two-part control. The HYPER toggle activates spatial diffusion mode and the HYPER depth knob sets the intensity. HYPER mode adds additional modulated delay taps to create an extremely dense, immersive chorus effect.

VOICES - Number of chorus voices: 2, 3 or 4. More voices produce a thicker ensemble.

SYNC - Locks the LFO rate to your DAW's tempo.

DIVISION - Beat division for tempo sync.

6.12 PHASER

A multi-stage allpass phaser that sweeps notch frequencies through the spectrum to create the characteristic swooshing, jet-like effect.

RATE - The speed of the internal LFO sweeping the allpass stages.

DEPTH - How far the notch frequencies sweep across the spectrum.

FEEDBACK - Controls the sharpness and intensity of the notches from -100% to +100%. Positive values create standard phaser resonance. Negative values invert the feedback creating a different tonal character with the peaks and notches swapped.

SPREAD - Notch spacing. Controls how the allpass stages are distributed across the frequency range from tightly clustered to widely spread.

STEREO - Stereo width decorrelation. Adds phase offset between the left and right channels for a wider spatial effect.

STAGES - The number of allpass stages: 4, 8 or 12. More stages create more notches in the spectrum for a denser, more complex phaser sound.

SYNC - Locks the LFO rate to your DAW's tempo.

DIVISION - Beat division for tempo sync.

Allpass coefficients are interpolated per-sample with a 3ms time constant to ensure click-free sweeps under modulation.

6.13 AUTO PAN

An automatic stereo panner that moves the signal between left and right channels using an internal LFO.

RATE - The speed of the panning movement.

DEPTH - How far the signal moves from centre. At 0% the signal stays centred. At 100% it sweeps fully from left to right.

SYNC - Locks the pan rate to your DAW's tempo.

DIVISION - Beat division for tempo sync.

6.14 BINAURAL

A binaural panner that positions the audio signal in True 3D space around the listener's head using interaural time and level differences. Designed for headphone listening.

AZIMUTH - The horizontal angle of the sound source. 0 degrees is directly in front, +90 is to the right, -90 is to the left and 180 is directly behind.

ELEVATION - The vertical angle. 0 degrees is at ear level, +90 is directly above and -40 is below.

WIDTH - Stereo width control. At 0% full binaural processing is applied. At 100% the original stereo signal passes through unprocessed. Values in between blend binaural positioning with the original stereo image.

6.15 STEREO WIDENER

A delay-based stereo widener using complementary comb filtering and Schroeder allpass decorrelation following the principles described by Lauridsen (1954). The technique splits the signal into low and high frequency bands using a Linkwitz-Riley 24 dB/oct crossover (4th-order) then applies different decorrelation delays to each band.

WIDTH - The overall stereo widening amount. At 0% the signal is unaffected. At 100% maximum widening is applied.

LO - Widens amount applied to low frequency band using a delay in the Haas range (~25ms).

HI - The widening amount applied to the high frequency band using a shorter delay (~15ms).

XOVER - The crossover frequency between the low and high bands.

SIDE - Toggles side-channel-only processing (M/S mode). When enabled only the side (stereo difference) component is processed leaving the mono centre untouched. This can produce a more natural widening effect on material with a strong centre like vocals or bass.

The visualiser shows a correlation meter indicating the stereo relationship of the output - from mono (fully correlated) through natural stereo to wide (decorrelated).

6.16 GRANULAR

A pure stereo granular processor that slices the incoming audio into tiny grains and reassembles them with independent control over density, size, pitch, direction, m/s spread and envelope shaping. This is one of NeuralDream's most parameter-rich engines offering deep control over the granular synthesis process.

6.16.1 WHAT SIZE ACTUALLY CONTAINS

Each size carries a pool of five grain lengths (in milliseconds) plus one base spawn interval. When a grain is spawned, its length is picked from the pool by slot index - every fifth grain gets the same length, so all five values cycle in rotation. The five values are prime-number sample counts, which prevents periodic phase reinforcement between simultaneous grains:

Size	Grain lengths (ms)	Base spawn interval (ms)
T (Tiny)	113 / 191 / 251 / 331 / 397	500
S (Small)	787 / 997 / 1103 / 1301 / 1511	2000
M (Medium)	1889 / 2111 / 2339 / 2543 / 2749	2500
L (Large)	3889 / 4111 / 4339 / 4561 / 4789	3000

Effect on grain length

Direct: S grains last ~0.8-1.5 s, M grains ~1.9-2.7 s, L grains ~3.9-4.8 s. That's a 3-5x difference between extremes. These are long grains by classic granular standards - NeuralDream is a wash/diffusion engine, not a micro-granular splatter engine and the lengths reflect that.

Effect on spawn frequency

This is the part that depends on which mode you're in.

SYNC mode on (straight = true): spawn rate is purely BPM-driven via DENSITY's musical divisions - 0% = 1 spawn per bar, 25% = 1 per half bar, 50% = 1 per quarter bar, 75% = 1 per eighth bar, 100% = 1 per sixteenth bar. The S/M/L size has no effect on spawn frequency in SYNC mode - only on grain length. So in SYNC, S/M/L is purely a length control.

SYNC mode off (free mode): spawn interval is exponentially interpolated between two anchors set by T/S/M/L:

- ✦ At DENSITY=0%, spawn interval = baseSpawnInterval from the table (T=500, S=2000ms, M=2500ms, L=3000ms). So S spawns more often than L at low density.
- ✦ At DENSITY=100%, spawn interval = averageGrainLength / maxGrains_ - fast enough to keep maxGrains_ slots saturated. Because S grains are shorter, they need replacing more often to maintain saturation, so S still spawns faster than L at high density (just for the inverse reason - short grains die faster).

So in free mode: smaller size = higher spawn rate at any DENSITY setting. Texturally this is "S = busy, percolating, dense cloud; L = sparse, big, slow, cinematic gestures" even at the same DENSITY.

PITCH - Pitch shift applied to each grain in semitones from -24 to +24.

SPRAY - How far back in the audio buffer grains can reach. Low values keep grains close to the current playback position for a tight, immediate sound. High values allow grains to draw from further back creating a smeared, time-stretched quality.

DENSITY - Grain spawn rate from sparse individual grains to dense overlapping clouds.

P.RAND - Random pitch variation per grain. Each grain gets a random pitch offset up to this amount creating a detuned, diffuse texture.

SMEAR - Per-grain pitch instability. Adds slow pitch drift within each grain's lifetime for an **organic**, tape-like quality.

PANORAMA - Acts as a M/S width control

SUP/COM - Sidechain dynamics control. Negative values duck the granular output when dry signal is loud. Positive values compress against the dry signal.

ATTACK - Grain envelope attack time from 0 to 5 seconds. At 0 seconds grains use a Hann window (smooth fade-in). Higher values create a slower fade-in for each grain.

RELEASE - Grain envelope release time from 0 to 5 seconds. At 0 seconds grains use a Hann window. Higher values create a slower fade-out giving grains a longer tail.

GRAINS - Maximum number of simultaneous grains from 1 to 24.

SIZE - Grain duration: T (tiny), S (small), M (medium) or L (large).

OCT - Toggles octave-shifted grains. When enabled additional grains are spawned one octave up for a shimmering harmonic quality.

5THS - Toggles fifth-shifted grains. When enabled additional grains are spawned a perfect fifth up adding a harmonic layer.

SYNC - Locks the grain density to your DAW's tempo.

FWD - increases the mix of forward playing grains in the final output. When this control is set to 0 there will be no forward grains created.

6.16.2 REV CONTROL

Behaviour at REV = 0% (default) there are no reverse grains spawned at all and therefore no extra buffer reads and no DSP cost.

Behaviour at REV > 0%: the parallel reverse stream spawns at the same density as the forward stream, runs through the same SPRAY/SMEAR/PITCH/PANORAMA/ATTACK/RELEASE settings but each grain walks the delay buffer backward at its assigned pitch (so reverse + Pitch=-12st = sub-octave reverse, exactly Submerge's "Sub" character; reverse + Pitch=0 = unison reverse, exactly Submerge's "Reverse" character). REV controls the relative level of the reverse stream against the forward stream - a true mix, not a coin-flip distribution.

Reference: Curtis Roads, Microsound (MIT Press, 2001), Chapter 5 - the canonical treatment of grain orientation and the perceptual character of reversed-envelope grains.

The visualiser shows individual grains as trapezoidal envelope shapes providing a direct visual representation of the grain activity.

FWD=100% / REV=100% - both streams play, mixed at full level. FWD=0% / REV=100% - forward dormant, reverse only. FWD=0% / REV=0% both pools dormant, output silent. Buffer keeps recording so streams resume cleanly when either knob is raised. Modulating FWD or REV - both pools' levels move smoothly. Spawn gate hysteresis at 0.001 keeps things tidy at low values.

6.16.3 THE GRANULAR VISUALISER



6.17 IR CONVOLVER

A convolution engine that applies the acoustic character of an impulse response (IR) to the incoming audio. This can simulate any acoustic space or equipment - concert halls, plate reverbs, speaker cabinets, vintage equipment and more.

The IR Convolver uses JUCE's partitioned FFT convolution for efficient real-time processing. An offline processing pipeline applies the following transformations to the loaded IR in sequence: offset, reverse, size, stretch, warp, decay and tilt with a pre-delay stage, capped at 30 seconds maximum length. Changes to these parameters are debounced (250ms) to avoid rebuilding the convolution on every knob movement.

LOAD - Opens a file browser to select a WAV or AIFF impulse response file.

RELOAD - Reloads the current IR file from disk. Useful if you have modified the file externally.

SAVE - Saves the processed IR (after all pipeline stages) as a new WAV file.

OFFSET - Trims the start of the IR, removing any initial silence or pre-ring.

REVERSE - Reverses the IR for reverse reverb effects.

SIZE - Scales the length of the IR.

STRETCH - Time-stretches the IR without changing its pitch.

WARP - Nonlinear time warping of the IR for creative reshaping.

DECAY - Controls the amplitude decay envelope applied to the IR.

TILT - A spectral tilt filter brightening or darkening the IR.

PRE - Pre-delay before the convolved signal, in milliseconds.

FEEDBACK - Feeds a portion of the convolved output back into the input for extended, evolving reverb tails. This is applied in real time with gain ramping for smooth adjustment.

NEURAL - Toggles neural IR generation mode. When enabled the LOAD/RELOAD buttons are replaced by SEED and DEPTH controls that generate procedural impulse responses using a temporal convolutional network (8 layers, kernel size 5, 8 channels, Kaiming initialisation, deterministic from seed). DEPTH maps linearly from 0.5 to 22 seconds.

Impulse responses are embedded in presets as FLAC-compressed base64 data ensuring complete portability. The visualiser shows a dual-layer waveform - a cyan ghost outline of the original loaded IR and a yellow live view of the processed IR after the pipeline.

6.17.1 ADDITIONAL INFORMATION ON LONG IRS

The IR Convolver will happily load impulse response files of any length but versions of NeuralDream prior to 1.1.40 struggled to play back IRs longer than around five seconds without hitting CPU overload. From v1.1.40 onwards three complementary techniques work together silently behind the scenes to make long, lush, expressive IRs genuinely usable in real-world sessions.

When an IR file is loaded, the engine first analyses its energy decay curve and trims any sub-audible tail content from the end before it reaches the convolver. This is perceptually transparent - the threshold sits 50 dB below the IR's own peak energy, well beneath the masking threshold of any program material - but on IRs with long quiet decay tails it can shorten the working buffer by 40-60 percent before any further processing happens. The Decay knob then operates as a true CPU control as well as a tonal one: pulling it down doesn't merely fade the tail, it physically truncates the IR at the point the decay envelope drops below an inaudible level, with a short cosine fade-out preventing any discontinuity. So lower Decay settings produce both the shorter reverberant tail you'd expect and a meaningful drop in convolver workload.

For impulse responses larger than approximately 2 MB (around 5.5 seconds of stereo audio at 48 kHz) a third technique engages automatically. The IR is split into three time segments - a short head section that runs at the full session sample rate, a mid section that runs at half the sample rate and a long tail section that runs at quarter rate. Because the upper-frequency content of any natural impulse response decays much faster than its lower-frequency content, this multi-rate scheme can deliver roughly a threefold reduction in CPU cost on long IRs without any audible loss of fidelity. The transition is automatic and silent; you'll see no UI change. Larger IRs that would previously have caused dropouts now load and play back smoothly, opening the door to the kind of long, beautiful convolver files that were always tempting but never practical.

6.18 PITCH SHIFTER

An STFT pitch shifter using Bernsee's smbPitchShift algorithm (WOL licence) with JUCE's vDSP-accelerated FFT on Apple Silicon. The engine uses a usual selectable from four FFT sizes - 512, 1024, 2048 or 4096 with 4x overlap. Larger sizes give better frequency resolution and stability at the cost of higher latency. Two independent pitch shift instances run in parallel for true stereo processing.

SHIFT - Pitch shift amount in semitones from -24 to +24 (two octaves in either direction). This is an integer-stepped control that snaps to semitone values.


FINE - Fine pitch adjustment in cents from -100 to +100 for precise detuning. Please note, when you apply modulation to this control (say to simulate vibrato), we specifically process the modulation source with a low-pass filter to smooth the modulation signal, which is optimised at between 3-6Hz.

SCALE - Selects a musical scale to constrain the pitch shift: Off, Major, Minor, Pentatonic Major, Pentatonic Minor, Blues, Major Triad or Minor Triad. When a scale is selected and CHORDS is off the SHIFT knob is quantised to valid scale degree intervals.

CHORDS - When enabled alongside a scale setting the engine performs per-bin spectral reshaping to produce chord harmonies from the selected scale rather than simple transposition.

FORMANT - Adjusting this control can sometimes improve the accuracy of the pitch-shifted audio, especially at shifts. The overhead is that 'behind the scenes', additional processing takes place which increases CPU load. However, if you leave the format control at zero this overhead will not be incurred.

The visualiser shows a flowing ribbon waveform with a chromatic pitch wheel indicating the current effective shift amount.



6.19 TAPE SIMULATION

A comprehensive analog tape emulation modelling the complete signal path of a tape recorder. The hysteresis stage uses the Jiles-Atherton equations solved with fourth-order Runge-Kutta (RK4) integration at 4x oversampling.

DRIVE - Input drive level controlling how hard the signal hits the virtual tape. Higher values push into saturation.

SATURATION - Controls the nonlinearity of the hysteresis model. At 0% the response is nearly linear. At 100% heavy tape compression and harmonic distortion are applied. Note that this control is inverted internally - higher values produce more saturation by reducing the Jiles-Atherton Ms and a parameters.

BIAS - The tape bias point. At 0% the bias is nominal producing a clean response. At 100% the tape is under-biased producing a warmer, more coloured tone with increased even harmonics.

SPEED - Simulates the effect of tape speed on high-frequency response. At 0% (slow speed) high frequencies roll off earlier simulating the head-loss characteristics of slow tape. At 100% the response extends to the full range.

WOW - Slow pitch modulation depth, simulating the gradual waver of tape transport rollers.

FLUTTER - Fast pitch modulation depth, simulating the higher-frequency shimmer that comes from capstan and pinch-roller imperfections.

NOISE - Adds tape hiss at an adjustable level.

NEURAL - A toggle on the top row of the panel that switches the hysteresis stage to a NAM neural network model. When enabled LOAD appears next to it for loading a .nam model file. The DRIVE, SATURATION and BIAS knobs are greyed out as they are replaced by the neural model's learned characteristics. DRIVE controls the input gain to the neural model. The rest of the tape path (loss filtering, head bump, wow/flutter, noise) continues to run as conventional DSP around the neural stage.

OUTPUT - Output level compensation.

The visualiser shows the hysteresis curve - the characteristic S-shaped relationship between input and magnetisation that defines the tape's distortion character.

6.20 NEURAL AMP

Loads and runs NAM (Neural Amp Modeller) model files for real-time amplifier, cabinet and pedal emulation. Uses the official NeuralAmpModelerCore C++ library by Steven Atkinson (MIT licence) supporting all four NAM architectures: WaveNet, LSTM, ConvNet and Linear.

LOAD - Opens a file browser to select a .nam model file. The model is loaded, prewired and ready to process audio. The model name is displayed on the panel.

INPUT - Input drive from -24 to +24 dB controlling how hard the signal hits the neural model.

OUTPUT - Output level from -24 to +24 dB.

TONE - A one-pole low-pass filter for darkening the output if the model is too bright.

The model file path is saved and restored with presets so the correct model is loaded automatically when recalling a preset.

The visualiser shows a VU-style peak meter with the loaded model name.

6.21 NEURAL FX

A randomised temporal convolutional network that generates unique nonlinear effects from a single seed number. Each seed produces a deterministic neural network with different learned characteristics - some produce subtle saturation while others create dramatic spectral reshaping, metallic textures or otherworldly distortions. The results are entirely unpredictable and exploring different seeds is part of the creative process.

SEED - An integer from 0 to 999 that deterministically generates the network's weights and architecture. The same seed always produces the same effect, so once you find a seed you like you can recall it exactly.

DEPTH - Network depth from 1 to 8 layers. More layers cascade more nonlinear transformations on top of one another, producing more extreme reshaping. Lower values keep more of the carrier audible.

RICHNESS - Channels per layer, selectable from 4, 8, 12, 16, 20 or 24. Higher values let the network represent denser, more harmonically rich effects at the cost of more CPU.

TEXTURE - Convolution kernel size, selectable from 3, 5, 7, 9, 11 or 13 samples. Larger kernels look further back in time at each layer, producing smoother, more temporal effects. Smaller kernels keep the response tight and immediate.

SPACE - Dilation factor from 1 to 4. Higher values stretch the receptive field of each layer out across more samples without increasing CPU, giving the engine a reverb-like or delay-like character.

TONE - Harmonic character selector with four positions: Warm (tube-like even-harmonic emphasis), Hard (clipping-style odd-harmonic emphasis), Bell (metallic resonant character) and Chaos (unpredictable, broken-machine textures).

The visualiser shows an animated mandala pattern that reflects the neural network's activity.

6.22 RAVE

A real-time audio variational autoencoder developed by Caillon and Esling (2021) at ACIDS, IRCAM, Paris. RAVE encodes incoming audio into a compact latent space (typically 8 dimensions) then decodes it back into audio that carries the timbral character of whatever dataset the model was trained on.

LOAD - Opens a file browser to select a .ts (TorchScript) model file.

BLEND - Crossfade between the carrier and modifier inputs through the latent space, from 0% (pure carrier) to 100% (pure modifier). Unlike a straight audio crossfade, this blend happens in the model's learned latent space, so intermediate positions produce timbres that the model considers naturally "between" the two - often closer to a continuous timbral morph than a layered mix.

DEPTH - How strongly the modifier signal influences the latent representation, from 0% (modifier has no effect, carrier passes through the model alone) to 100% (modifier exerts full influence). Use lower values to keep the carrier's identity dominant while only colouring it with the modifier; higher values let the modifier reshape the carrier more aggressively.

DIMS - Selects how many of the latent space's dimensions the modifier controls, from 1 to 8. With Dims at 1 only the most prominent latent axis is affected - typically corresponding to pitch or overall brightness depending on how the model was trained. As Dims rises, more of the timbral structure comes under modifier influence, until at 8 the full latent representation is being shaped.

TEMP - Latent-space temperature from 0% (deterministic, fully governed by the input) to 100% (wild - substantial random variation injected into the latent vectors). Low values give predictable behaviour; high values produce ever-shifting timbral fluctuations even when the input is steady.

The RAVE engine operates as a dual source engine. The primary audio is encoded and decoded through the model while the modifier input can influence the latent representation.

RAVE requires the ML-enabled build of NeuralDream (built with LibTorch). The engine is registered in the menu only when ML support is compiled in.

The visualiser shows streaming latent space values as flowing coloured traces - one per latent dimension.

6.23 SPECTRAL MORPH

A dual source spectral processor that blends the spectral characteristics of two audio signals. The carrier and modifier signals are analysed using FFT (4096 points, Hann window, 75% overlap) and the output spectrum is a weighted blend between them.

MORPH - Controls the blend between carrier and modifier spectra. At 0% you hear the carrier unchanged. At 100% the output takes on the spectral shape of the modifier while retaining the carrier's phase structure.

SHIFT - Shifts the modifier spectrum before morphing in semitones from -24 to +24. This allows you to transpose the modifier's spectral contribution relative to the carrier for creative harmonic effects.

SMOOTH - Temporal smoothing of the spectral analysis. Higher values create a slower, more gradual morph that smears transients.

SIDE - Side channel level. At 0% the output is pure morphed mono. At 100% the original stereo information passes through.

The engine uses mid-side stereo processing internally for phase-coherent results.

6.24 SPECTRAL GATE

A dual source spectral processor that uses the modifier signal's spectrum as a frequency-dependent gate for the carrier signal. Frequency bins where the modifier has energy are allowed through while bins where it does not are attenuated. This creates an effect where the carrier signal is sculpted by the spectral shape of the modifier - only the frequencies present in both signals survive.

THRESHOLD - Controls how much modifier energy is needed at each frequency bin to open the gate. Low values let more of the carrier through. High values require strong modifier energy to open each band.

BANDS - Number of frequency bands from 4 to 128. Low values (4-8) produce broad dramatic gating. Medium values (32) give a classic vocoder-like quality. High values (128) provide detailed frequency-by-frequency gating.

RANGE - Floor gain for closed bands. At 0% closed bands are fully silenced (hard gate). Higher values let some carrier signal bleed through even in closed bands for a softer, more subtle effect.

ATTACK - Gate opening speed from 0.1 to 50 ms. Fast settings preserve transients. Slow settings create swelling effects as bands open gradually.

RELEASE - Gate closing speed from 1 to 500 ms. Short values give tight staccato gating. Long values give flowing sustained tails.

SIDE - Side channel level. At 0% the output is pure gated mono. At 100% full stereo passthrough.

6.25 VOCODER

A dual source spectral vocoder that imposes the spectral envelope of the modifier signal onto the carrier signal. The classic application is feeding a synthesiser as the carrier and a voice as the modifier to create the "talking synthesiser" effect.

BANDS - Number of vocoder bands from 4 to 128. Low values (8) produce a thick robotic quality. Medium values (32) give the classic vocoder sound. High values (128) provide hi-definition vocoding that preserves more of the modifier's nuance.

FORMANT - Shifts the modifier's spectral envelope in semitones from -24 to +24. This changes the perceived pitch of the formant structure without changing the carrier's pitch - effectively raising or lowering the "voice" of the vocoder. Negative values make formants sound deeper while positive values make them brighter and thinner.

ATTACK - Envelope follower attack time from 0.1 to 50 ms. Controls how fast each band responds to rising energy in the modifier signal.

RELEASE - Envelope follower release time from 1 to 500 ms. Controls how fast each band responds to falling energy. Shorter values give a tighter, more articulated vocoder. Longer values smooth out the response.

SIDE - Side channel level. At 0% the output is pure vocoded mono. At 100% full stereo passthrough.

The vocoder uses the same FFT framework as the other spectral engines (4096 points, Hann window, 75% overlap, mid-side stereo).

6.26 LIVE CONVOLVE

A dual source engine that performs real-time convolution between two live audio signals. Unlike the IR Convolver which uses a static impulse response file Live Convolve treats the modifier signal as a continuously updating impulse response creating evolving, reactive convolution effects.

DEPTH - Controls the depth of the convolution effect. At 0% you hear the carrier unchanged. At 100% full convolution is applied.

TILT - Spectral tilt compensation from -24 to +24 dB. Convolution naturally darkens the output because multiplying two spectra together tends to attenuate high frequencies. The tilt control lets you compensate for this by boosting or cutting the high-frequency balance.

SMOOTH - Temporal smoothing of the modifier spectrum. Higher values create a more gradual, evolving convolution rather than reacting instantly to every change in the modifier.

SIDE - Side channel level. At 0% the output is pure convolved mono. At 100% full stereo passthrough.

6.27 RING MODULATOR

A dual source ring modulator that multiplies the carrier and modifier signals together sample by sample. Ring modulation produces sum and difference frequencies creating inharmonic, metallic and bell-like timbres.

DEPTH - Controls the blend between the dry carrier and the ring-modulated output. At 0% you hear the carrier unchanged. At 100% you hear the fully ring-modulated signal.

The modifier signal is whatever audio source is assigned to the engine's modifier slot - another layer's output, a sample or the plugin input. The character of the ring modulation depends entirely on the spectral relationship between the carrier and modifier signals.

6.28 SPECTRAL STYLE TRANSFER

A dual source engine that performs cross-synthesis by transferring the spectral envelope of the modifier signal onto the carrier signal. Based on the spectral envelope transfer technique described by J.O. Smith III at CCRMA, Stanford (2020).

The algorithm works by extracting spectral envelopes from both signals using frequency-domain running-average smoothing of log-magnitudes, computing per-bin gain differences between the two envelopes and applying those gains to the carrier spectrum while preserving its phase. The result takes on the tonal shape of the modifier while retaining the temporal and pitch characteristics of the carrier.

DEPTH - Controls the strength of the envelope transfer from 0% (no effect) to 100% (full transfer).

SMOOTH - The cepstral smoothing order (5-80) controlling how detailed the spectral envelope extraction is. Lower values capture fine spectral detail while higher values extract only the broad spectral shape.

SIDE - Controls the modifier signal balance.

LO / HI - Frequency range limits for the envelope transfer allowing you to restrict the effect to a specific frequency band.

FREEZE - Captures the current modifier spectral envelope and holds it, allowing you to continue processing the carrier against a frozen snapshot of the modifier's spectrum even after the modifier audio stops.

This engine uses FFT 2048 with Hann window and 75% overlap (512-sample hop), introducing 1536 samples of latency - about 32 ms at 48 kHz or 35 ms at 44.1 kHz. Latency is reported to the host so plugin delay compensation handles alignment automatically.

The visualiser shows dual waveforms for the carrier and modifier signals.

6.29 OCTAVE FILTER

A 10-band stepped octave band-pass filter, inspired by the classic portable octave analysers of the 1960s such as the Brüel & Kjær 1613 portable octave analyser, a precision passive LC instrument originally designed for acoustic measurement work in broadcast, research and live sound. The 1613 found unexpected creative use in the BBC Radiophonic Workshop and European electronic music studios where its steep band isolation and characterful resonance became tools for spectral sound design.

NeuralDream's implementation uses three cascaded Cytomic ZDF SVF (zero-delay feedback state variable filter) stages per band giving 6-pole (36 dB/octave) slopes that match the tight band isolation of the original hardware. The ten bands are centred on the ISO standard octave frequencies: 31.5, 63, 125, 250, 500, 1K, 2K, 4K, 8K and 16K Hz.

Two operating modes are available:

Sweep mode - The FREQ knob selects a single band with smooth crossfading between adjacent bands as the knob moves between positions. Detent positions are provided at each of the ten band centres for precise selection. The FREQ knob displays the current frequency in Hz and is fully modulatable for swept filter effects. Weight interpolation runs per-sample to eliminate zipper noise during sweeps making this control ideal for LFO and envelope follower modulation.

Bank mode - All ten bands are active simultaneously with independent level controls for each. This turns the engine into a 10-band graphic equaliser-style filter bank where you can sculpt the spectral balance by raising or lowering individual octave bands.

The MODE button toggles between Sweep and Bank modes.

Controls are as follows:

FREQ - Band selector frequency. In Sweep mode this selects which octave band is active with smooth crossfading between adjacent bands. Detent positions snap to each of the ten ISO centre frequencies. Displays the current frequency in Hz. In Bank mode this control has no effect.

Q - Filter resonance (impedance). Controls the sharpness of the bandpass peak. At minimum (0%) Q is set to 1.0 giving clean octave-bandwidth filtering with unity passband gain and no resonance. Increasing Q narrows the passband and adds a resonant peak at the centre frequency. This recreates the idea behind the impedance control on analysers of that kind, where impedance mismatch directly controlled the filter damping and therefore the resonant character.

DRIVE - Input saturation applied before filtering. Uses tanh soft clipping to add warmth and harmonic content to the signal before it reaches the bandpass stages.

CLIP - Output soft clipping applied after filtering. Tames peaks and adds character to the filtered output.

Band level knobs (Bank mode) -- Ten individual level controls, one per octave band. Each knob sets that band's contribution from 0% (silent) to 100% (full level). These are labelled 31, 63, 125, 250, 500, 1K, 2K, 4K, 8K and 16K.

The visualiser shows the ten octave bands as vertical columns with the currently active band highlighted in the engine's warm amber accent colour. The column height represents the weight of each band in the current configuration.


6.30 LEVEL/PAN UTILITY

A simple two-knob utility engine: an automatable output level control and a stereo pan. The purpose of this engine is pragmatic - the layer-level volume and pan controls at the edge of each layer strip cannot currently be automated from the host DAW and some users need host automation over these quantities. Dropping a Level/Pan Utility into any slot gives you two parameters the DAW can see, record and automate freely, while leaving the layer-level controls available for manual, non-automated mixing.

The LEVEL knob runs from 0 to 100% using a perceptually-linear mixing-fader taper. Fully anticlockwise at 0% is silence. Unity gain (0 dB) sits at approximately 79% - the equivalent of a channel fader at its 0 dB marking - so the bulk of the knob's travel covers the useful working range below unity. The final 21% of travel opens up a boost region that reaches +12 dB at fully clockwise, giving headroom for layers that need emphasis. This is the standard fader curve used by most professional DAW channel strips.

The PAN knob uses a -3 dB equal-power law: fully anticlockwise (-100) sends the signal hard left, centre (0) is dead centre with both channels at full gain and fully clockwise (+100) sends hard right. At the centre position the law introduces no loss; as the knob moves toward either extreme, the opposite channel attenuates smoothly while the near channel holds its level, keeping perceived loudness roughly constant across the stereo field.

Both parameters are smoothed internally with a ~5 ms one-pole filter, which prevents zippering under rapid DAW automation without introducing any audible lag. The engine is zero-latency and has no other controls - it does exactly what its name suggests and nothing more, which is the whole point.



6.31 CROSS-SYNTHESIS

A dual source spectral processor that fuses the timbres of two audio signals by separately morphing their spectral envelopes and their spectral detail. Unlike Spectral Morph which blends complete spectra and Spectral Style Transfer which copies one envelope onto the other Cross-Synthesis decomposes each signal into a smooth formant shape (the envelope) and a fine harmonic structure (the detail) and interpolates each component independently. The result is a true timbral hybrid - the morph can feel like a continuous slide between two distinct instruments rather than a layering of one over the other.

The algorithm follows the cross-synthesis formulation described by J.O. Smith at CCRMA, Stanford (Spectral Audio Signal Processing, 2011). For each FFT frame the carrier and modifier magnitudes are split into envelope and detail using cepstral liftering: the log-magnitude is transformed to the cepstrum, low-frequency coefficients are retained to recover the broadband shape and the residual gives the harmonic detail. The two envelopes and two detail spectra are then geometrically interpolated in log space - equivalent to a perceptually linear blend of formant peaks rather than the destructive sum produced by linear-amplitude morphing. The recombined magnitude is recoupled with the carrier's phase before inverse FFT and overlap-add reconstruction.

MODE - Cross-synthesis algorithm selector with four positions:

- ◆ Morph (default) - Blends the carrier and modifier spectra continuously, using the cepstral envelope-and-detail decomposition described in the algorithm overview above. The MORPH knob controls the blend position. This is the smooth-fade-between-two-instruments mode.
- ◆ Swap - Takes the magnitude spectrum from the modifier and pairs it with the phase spectrum of the carrier. The result has the timbre of the modifier but the time-domain articulation of the carrier. Useful for imposing one instrument's tonal character onto another's rhythmic envelope.
- ◆ iSwap - The mirror of Swap. Takes magnitude from the carrier and phase from the modifier. Often produces a more dramatic, less predictable hybrid than Swap.
- ◆ Multiply - Multiplies the two magnitude spectra together bin by bin, retaining the carrier's phase. The result emphasises frequencies that are present in *both* signals and silences those present in only one. Particularly effective for spectral intersection effects and for processing tonal material against a percussive modifier.

MORPH - The MORPH knob's behaviour adapts to the mode: in Morph mode it controls the envelope-blend position, while in Swap, iSwap and Multiply modes it controls the wet/dry balance between the chosen algorithm's output and the pure carrier.

TILT - Post-morph spectral tilt from -12 to +12 dB per octave referenced to 1 kHz. Negative values darken the output, positive values brighten it. Useful to compensate when one source is much brighter than the other or to shape the morphed result tonally.

FORMANT - Shifts the modifier's spectral envelope along the frequency axis from -12 to +12 semitones independently of the modifier's pitch. Negative values lower the formant region producing a deeper, warmer timbre. Positive values raise it producing a brighter, thinner timbre. The carrier's pitch is unaffected.

FREEZE - Latches the current modifier envelope. While engaged the held envelope continues to **shape** the morph even when the modifier audio falls silent. Combine with DENSITY at low values to capture a particular formant snapshot from a transient moment in the modifier.

DENSITY - Controls how quickly the modifier envelope follows the live modifier signal from 100% (instantaneous tracking, every frame fully refreshed) down to 0% (a six-second time constant across modifier changes). At low values the envelope drifts slowly across modifier variations producing a smeared, time-averaged formant character useful for evolving pad textures. At high values the morph follows the modifier in real time. The control uses continuous one-pole smoothing rather than discrete snapshots so transitions between settings are always musical.

MOD GAIN - Modifier input gain from 0% to 200%. Provides headroom adjustment when the modifier source is very quiet or very loud relative to the carrier.

QUALITY - Cycles through FFT window sizes from 512 to 4096 samples. Larger windows give better frequency resolution and a more stable envelope extraction at the cost of higher latency and slower transient response. Smaller windows tighten the time response and reduce latency but coarsen the envelope. The default is 4096 which suits sustained material.

The engine introduces latency equal to the selected FFT window plus a short internal lookahead - approximately 98 ms at QUALITY 4096 / 44.1 kHz, dropping to 17 ms at QUALITY 512. Latency is reported to the host so transport position remains sample-accurate.

Cross-Synthesis uses mid-side stereo processing internally for phase-coherent results.

The visualiser shows dual waveforms for the carrier and modifier signals.

6.32 BITFOLD

A three-mode digital destruction engine. The MODE button cycles between FOLD (analytical sine wavefolder), CRUSH (continuous bit-depth reduction with TPDF dither and noise shaping) and REDUCE (sample-rate reduction with smoothing and clock jitter). The first three knobs swap function with the mode; TONE is universal. An internal +15 dB lookahead limiter prevents extreme settings from running away.

6.32.1 FOLD MODE - DRIVE, STAGES, SYMMETRY

DRIVE - Crossfaded fold amount. At 0% the engine is bypassed (output equals input). Higher values push the signal through the fold curve producing rich harmonic content.

STAGES - Number of cascaded fold stages from 1 to 4 with 0.7x inter-stage attenuation. Each stage produces qualitatively different timbres. Four stages give the dense, complex folding associated with classic West Coast oscillators.

SYMMETRY - DC offset added before the fold. At 0% the fold is symmetric and produces pure odd harmonics. Higher values introduce even harmonics for tube-like warmth.

The visualiser draws the live transfer curve in glowing amber with a hot-red dot tracing the current sample's position. As you turn DRIVE, STAGES or SYMMETRY the curve updates in real time.

6.32.2 CRUSH MODE - BITS, DITHER, NOISE

BITS - Logarithmic bit-depth control. Knob 0% is 16 bits (transparent), 25% is 8 bits, 50% is 4 bits, 75% is 2 bits, 100% is 1 bit (square wave). Continuous fractional values work fine.

DITHER - TPDF (Triangular Probability Density Function) dither added before quantization. At 0% pure truncation produces signal-correlated harmonic distortion (the classic SP-1200 character). Higher values convert that distortion into noise.

NOISE - Three-position noise shaping: off, 1st-order, 2nd-order. Pushes quantization noise spectrum into less-audible high frequencies.

The visualiser draws the audio output as glowing green dots scrolling across a horizontal cyan grid of quantization levels. As BITS drops, the grid becomes more visible and the dots cluster onto the level lines.

6.32.3 REDUCE MODE - RATE, SMOOTH, JITTER

RATE - Effective sample rate, log-mapped from 100 Hz at 0% knob to your project's actual sample rate at 100% knob (transparent).

SMOOTH - Crossfade between sample-and-hold (0%, hard staircase) and cubic Hermite interpolation (100%, sloped curves between hold-points).

JITTER - Random variation in the sample clock. Adds the warble and instability of vintage hardware DACs.

The visualiser overlays the smooth input as a thin sky-blue line and the stepped output as a brighter indigo line. At low RATE the output line shows visible plateaus; SMOOTH softens them; JITTER makes plateau widths irregular.

TONE - Universal output low-pass filter, log-mapped from 200 Hz to 20 kHz. Defaults fully open. Useful for taming the harshness of any of the three modes.

6.33 LUSTRE

A bright modern algorithmic reverb built on an eight-line feedback delay network with Walsh-Hadamard mixing. The topology follows the lossless prototype FDN of Stautner and Puckette (1982) and Jot (1991), with per-line one-pole damping, per-line LFO modulation for organic motion, Schroeder allpass pre-diffusion for early reflection density, M/S stereo processing, DC blocking and an output high-shelf for modern sheen.

Two independent reverb engines run in parallel, one per channel, with a counter-phased stereo wander LFO offsetting their delay lines so the L and R fields breathe slightly out of step. This produces a true stereo image without any cross-channel mixing - the left tail is genuinely the left tail, not a derived widening of a mono reverb.

PRE-D - Pre-delay before reverb onset, from 0 to 500 ms. Useful for preserving transients ahead of a long tail. When SYNC is engaged the pre-delay locks to your DAW tempo and the BeatDiv combo selects the beat division.

SIZE - Plate area, scaling the perceived reverb space from a small chamber to a vast cathedral.

DECAY - Reverb tail length. Drives the FDN feedback gain to set the RT60 from short and percussive to long and immersive.

DAMP - High-frequency damping in the feedback path. At 0% the damping filter sits above 20 kHz which is effectively transparent - bright by default. At 100% the cutoff drops to roughly 4 kHz for a warm tail. The damping is in the feedback loop, so brighter material darkens generation by generation.

WIDTH - Stereo width via M/S processing applied at the output. 0% collapses to mono, 100% spreads the image to the panel edges, values above 50% widen beyond the natural FDN spread.

HI PASS - Input high-pass cutoff from 20 Hz - 1 kHz. Tames bass buildup as it enters the reverb.

LO PASS - Input low-pass cutoff from 1 kHz to 22 kHz. Darkens the input ahead of the reverb so high-frequency content does not enter the network at all.

DENSITY - Pre-diffusion strength, controlling the density of early reflections. Low values let transients survive into the tail. Higher values smear the input more aggressively before it hits the FDN.

MOD - Delay-line modulation depth. Each FDN line carries its own slow LFO that modulates its delay length by a few samples. At 0% the reverb is glassy and static. As MOD rises the tail develops a gentle breathing motion that prevents standing-wave artifacts and adds an organic quality.

MOD RT - Controls the speed of the per-line LFOs from a slow drift through to a faster shimmer.

BOOST - Two-band "Sparkle" EQ applied at the engine output, scaling together from 0% (both bands flat) to 100%. The first band is a presence bell at 3.5 kHz with $Q \approx 0.9$, reaching up to about +8 dB at full setting and targeting the Fletcher-Munson sensitivity peak where the ear hears most distinctly. The second band is an air high-shelf at 10 kHz reaching up to about +5 dB at full setting, adding the open top-end shimmer that defines this engine. The default of 30% gives the modern "sheen" character that distinguishes Lustre from a more classical FDN response. Pull it down for a flatter, more vintage tail; push it up for aggressive air.

SHIMMER - Feeds pitch-shifted copies of the reverb back into the tank so the tail climbs. At 0 percent the shimmer path is off and costs nothing; every existing Lustre preset sounds exactly as it did. As the amount rises the shifted voices grow louder in the tail. Two shifted voices are used, an octave and a fifth by default and their intervals are set by SHIFT 1 and SHIFT 2 below.

MODE - Where the shimmer takes its signal from. INPUT shimmers only the incoming audio, so the shifted voices are heard once and do not regenerate. TANK shimmers the reverb itself, so each pass through the tank climbs again and the tail ascends. BOTH does both at once: the input is shimmered on the way in and the tank keeps climbing.

SHIFT 1 and SHIFT 2 - The interval of each shimmer voice in semitones, from two octaves down to two octaves up. The defaults are +12 (an octave) and +19 (an octave and a fifth). Negative values give a descending shimmer.

WANDER - A slow random drift on the shimmer's pitch, in cents, so the shifted voices are not surgically exact. A small fixed drift has always been part of the shimmer; WANDER adds a random walk on top of it, different on the left and right, up to about plus or minus fifteen cents at 100 percent. That is well inside where an octave still reads as an octave but it stops the shimmer from sounding like a machine. At 0 percent nothing is added.

DRIFT - Random motion on every one of the eight delay lines, sitting underneath MOD. MOD is a sine wave and past a small amount the ear finds its cycle and hears vibrato. DRIFT is a slow random walk on each line, each line on its own and the left and right sides differently, so the tail keeps moving with no cycle to find. It is the control to reach for when the tail feels static but more MOD makes it wobble. At 0 percent nothing is added; the reverb is exactly as it was.

BLOOM - Makes the reverb swell in after a rest rather than arriving all at once. With BLOOM up, the first sound after a moment of silence enters the reverb quietly and opens over a period set by the knob, up to about a second and a half at 100 percent. Play continuously and nothing changes; leave a gap and the next phrase blooms. At 0 percent the input arrives instantly as before.

FREEZE - Holds whatever is in the reverb indefinitely. When you switch it on, the current contents of the tank are captured into a second reverb that recirculates without loss and without damping, so the held sound sustains for as long as FREEZE stays on and keeps breathing under MOD and DRIFT. The live reverb carries on underneath, so you can keep playing over the frozen sound and every new note still gets its own reverb. If shimmer is on, the frozen sound keeps climbing. Switching FREEZE off lets the held sound fade away over a quarter of a second rather than cutting it dead. FREEZE is off in every preset unless you save it on and the V button never switches it on.

SYNC + DIVISION - Tempo-lock the pre-delay to your DAW. When SYNC is on the PRE-D knob is hidden and the DIVISION combo selects from 1/32 up to 4 bars. The SYNC button and the DIVISION combo sit together at the bottom left of the panel, directly under HI PASS, with the DIVISION combo appearing just below the SYNC button when SYNC is on. Useful when you want the pre-delay to align with the song's groove rather than sit at an arbitrary millisecond value.

The visualiser is one of the most distinctive in NeuralDream - a 3D banded reverb-tail rendering. Ten frequency bands from 60 Hz to 18 kHz are tracked in a 4-second history buffer, each band rendered as a receding ribbon stretching from the front of the panel up toward a vanishing point at the top. Each band is coloured along the spectrum from deep blue at the lowest frequency to bright pink- red at the highest, with per-band brushes tuned so high-frequency bands receive a softer vertical falloff to keep their distant slices seamless. A pale blue haze halo marks the vanishing point. The full panel height is filled with the receding tail.

6.34 DREAM DELAY

A three-tap tempo-locked tape echo inspired by the delay in the u-he Diva synthesiser. Three independent fractional delay lines (Left, Centre, Right) each lock to integer thirty-second-note multiples, sweepable from 1/32 (one thirty-second) up to 128/32 (four bars at 4/4). Each tap-time knob travels continuously under your cursor for smooth automation but reports an integer value, with magnetic detents at the musically meaningful boundaries (1/32, 1/16, 1/8, 1/4, 1/2, whole note, dotted whole, two bars, three bars, four bars).

Centre defines the feedback period - the centre tap's output is the source of a single shared feedback bus that runs through high- pass and low-pass filters in series before being summed back into all three delay-line inputs alongside the new audio. With HP at minimum and LP at maximum the cascade is unity-gain in the audio band, so 100% feedback produces a true infinite repeat; pulling either filter inward bleeds energy out of the loop generation by generation.

L TIME / C TIME / R TIME - The three tap times in thirty-second- note steps. The Left and Right taps fan out to their respective output channels for stereo width (the ping in the field) while the Centre tap is mono-positioned and goes equally to both channels. At identical settings on all three knobs every repeat lands at the same moment, which gives you a simple stereo delay; spreading them apart produces a true three-tap pattern.

CENTRE VOL - Level of the centre tap in the output mix. Drives how loud the mono-positioned repeat sits relative to the stereo pair.

SIDE VOL - Level of the Left and Right taps. Controls the strength of the stereo ping. With Centre Vol at zero and Side Vol up you get a ping-pong delay; with Side Vol at zero you get a mono delay.

WOW - Tape wobble. At 0% the modulation LFO is bypassed entirely. Above zero the knob sweeps the LFO rate from a subtle drift around 0.3 Hz at low values up to a seasick wobble around 6 Hz at the top. The depth is fixed at a musical amount (about ± 3 ms peak-to-peak). A single shared LFO modulates all three taps in unison, so the wobble is coherent across the field - no phasey chorusing between taps.

FEEDBACK - Regeneration amount. With HP at minimum and LP at maximum, 100% produces an infinite repeat as documented; with either filter brought in, the maximum sustainable feedback drops accordingly.

HP - High-pass cutoff in the feedback path, from 20 Hz to 2 kHz. Removes low-frequency buildup from successive repeats.

LP - Low-pass cutoff in the feedback path, from 200 Hz to 18 kHz. Darkens successive repeats for a natural tape-style tail.

L SHIFT / C SHIFT / R SHIFT - Independent pitch shifters per tap, each ranging from -12 to +12 semitones in integer-semitone steps with a centre detent at zero. Each tap's output is shifted by its own knob amount using the same Bernsee STFT phase vocoder as the standalone Pitch Shifter engine. Use them to turn straight echoes into octave-up shimmer, harmonised intervals or descending pitch trails.

IN LOOP - Routes the centre tap's pitch-shifted version back into the feedback bus, so each successive repeat accumulates a further C SHIFT amount. With C SHIFT at +12 and IN LOOP on you get the classic studio shimmer effect - first echo at +12, second at +24, third at +36, etc, climbing into the air until the feedback fades out. With IN LOOP off the shifters apply only at the output and every repeat sits at the same pitch.

The visualiser is a three-lane cloud-particle field. The Left lane sits at 25% of the panel width, Centre at 50%, Right at 75%. Each tap fires periodically in lockstep with its own delay time - every L TIME samples the Left lane spawns a cloud, every C TIME samples the Centre lane spawns a cloud and so on. Cloud size is driven by the tap's recent envelope at the moment of firing, so the visible repeats shrink as feedback decays. The Wow LFO's phase is shared between audio and visuals, so increasing Wow visibly makes the entire field seasick in step with the audible wobble.



When pitch shift is active each cloud carries a tint reflecting its tap's shift amount captured at spawn time. Positive shifts slide the cloud halo and core toward pale blue (BLUE-SHIFTED, like a star moving toward the observer); negative shifts slide toward pink (RED-SHIFTED). The strength of the tint scales with the shift magnitude up to about ± 9 semitones where it saturates fully. Each cloud locks its tint at spawn so a moment-by-moment record of the pitch trajectory is preserved as the clouds drift and fade - useful visual feedback when working with IN LOOP cascading shimmer.

Dream Delay is true stereo from input to output. Your left channel and your right channel travel through their own delay lines, their own feedback filters and their own pitch shifters and are never summed together on the way. The left repeat carries your left channel, the right repeat carries your right and the centre repeat carries both. A wide stereo source therefore keeps its width and its top end when you turn the mix fully wet, rather than collapsing.

6.35 APERTURE

A brickwall STFT band gate that passes audio inside an adjustable frequency window and silences everything outside it. The engine builds a frequency-domain mask in real time and applies it bin-by-bin to a sliding Short-Time Fourier Transform of the signal, following the overlap-add framework described by Allen and Rabiner (1977) in "A Unified Approach to Short-Time Fourier Analysis and Synthesis" published in the Proceedings of the IEEE and developed further in J.O. Smith III's STFT material at CCRMA, Stanford. The implementation uses a 2048-point FFT with a 512-sample hop (75% overlap, four-times overlap factor), Hann windowing applied at both the analysis and synthesis ends to ensure constant overlap-add reconstruction and a fixed COLA gain of $2/3$. Latency is 1536 samples - about 32 ms at 48 kHz - and is reported to the host so DAW plugin delay compensation handles alignment automatically.

Unlike a cascade of LP and HP filters, Aperture has no phase distortion across the pass band and an arbitrarily steep brickwall response when SLOPE is at minimum. It is useful for surgical band isolation, removing wide swathes of spectrum that LP/HP cascades cannot cleanly cut and - with INVERT engaged - for notching out targeted frequency regions. Because it operates in the STFT domain transients exhibit a small amount of pre-ring of roughly half the FFT size so very percussive material may benefit from following Aperture with a transient-preserving stage.

LO - The low edge of the pass band from 20 Hz to 20 kHz. Frequencies below this edge are attenuated according to the SLOPE setting.

HI - The high edge of the pass band from 20 Hz to 20 kHz. Frequencies above this edge are attenuated according to the SLOPE setting.

SWEEP - Slides the whole band up or down the spectrum without changing its width, from minus 100 to plus 100 percent with zero at the centre. LO and HI still set where the band sits at rest; SWEEP moves both edges together from there, upwards for positive values and downwards for negative ones, each by the same amount along the same curve the LO and HI knobs use, so a band that is an octave wide stays an octave wide as it travels. The knob sits alone on the top row of the panel above the others. It is the control to modulate when you want the band to move: an envelope on SWEEP makes the band jump up with every hit and drift back and a slow LFO on SWEEP with a faster one on LO and HI together gives a band that wanders across the spectrum

while breathing in and out. When a sweep pushes one edge to the end of the range, that edge stops there while the other keeps moving, so at the extremes the band narrows and eventually closes, which can be used deliberately or avoided by keeping the modulation depth short of the ends. The visualiser draws the band where it is after the sweep, so the window you see is always the one you hear.

SLOPE - Edge taper from 0 to 100%. At 0% the band edges are perfectly brickwall - energy outside the LO-HI window is silenced completely with no transition region. As SLOPE rises the edges soften into a smooth roll-off, reaching a gentle bell-like shape at 100%. Modulating SLOPE lets you sweep between surgical isolation and gentle band emphasis.

TILT - Spectral tilt across the pass band from -6 dB to +6 dB. The pivot is the geometric mean of LO and HI: at this centre frequency the gain is 0 dB, at HI the gain reaches +TILT and at LO it reaches -TILT, scaling logarithmically in between. Positive values brighten the pass band, negative values darken it, in either case without shifting the perceived band centre.

INVERT - Flips the mask. When engaged, the LO-HI window becomes a notch instead of a pass band - energy inside the window is silenced and everything outside is preserved. The other controls continue to behave as expected, just with their effect inverted.

The visualiser shows the spectrum of the input signal as a log-frequency display from 20 Hz to 20 kHz with grid lines and labels at 20, 50, 100, 200, 500 Hz and 1, 2, 5, 10 and 20 kHz. A live frequency mask drawn from the current LO, HI, SLOPE, TILT and INVERT values shows the gate shape in cyan for the pass region and yellow for the blocked region. Two cyan hairlines mark the LO and HI edges, after SWEEP, so you can always see exactly where the band sits even when SLOPE is wide and the band shape is gentle. The visualiser uses an independent 1024-point FFT separate from the DSP path so the display stays smooth and responsive without affecting the audio processing.

6.36 RESONANCE FIELD

A spectral resonator that turns any incoming audio into a tunable harmonic field. Rather than filtering the input or adding voiced oscillation on top of it, Resonance Field extracts the input's spectral energy frame by frame and feeds it into a bank of tuned spectral resonators - one per harmonic of one or more fundamental pitches - where each resonator stores energy and releases it back into the output over time. The result is a sustained pitched ringing that takes its character from your input but its musical content from the voicing knobs.

The engine builds on the same Short-Time Fourier Transform overlap-add framework as Aperture (Allen and Rabiner 1977, J.O. Smith III at Stanford CCRMA). A sliding STFT analyses the input at the user-selected block size and hop factor; the analysis magnitude is reshaped by a harmonic mask, fed through a per-bin Spectral Mass-Damper accumulator whose decay coefficient controls how long energy persists in each bin, recombined with the input phase and reconstructed via overlap-add. The brain of the engine is the per-bin SMD: a first-order recursive accumulator that gives every spectral bin its own controllable ring time, producing the long sustained resonator tail without any time-domain delay-line resonator and without the energy buildup or feedback howling that traditional resonator structures suffer from.

Resonance Field works on any input. Fed a drum loop it produces a sustained pitched wash whose attack envelope and rhythmic accents come from the loop and whose pitch content comes from the voicing knobs. Fed a pad it produces a denser harmonised version of the pad's spectrum that sustains beyond the original notes. Fed speech it produces a slowly evolving choir-like texture. The denser and more transient your input the more dramatic the transformation - quiet sustained material may need the GATE control engaged to keep low-level noise out of the resonator. Two voicing modes are available. Internal mode (the default) is driven by the four V1-V4 knobs on the panel - V1 sets the fundamental MIDI note number and V2 / V3 / V4 add optional further voices at user-set semitone offsets from V1. MIDI mode replaces those four voices with notes received from the host: up to twelve simultaneous polyphonic voices, with FIFO voice stealing when the chosen polyphony cap is exceeded. In either mode the resonator excites the same way from your input audio - the voicing controls only decide which pitches the harmonic mask is tuned to.

The panel is laid out as four blocks reading left to right in the order the signal flows, separated by fine vertical lines. The VOICING block on the left sets which pitches the resonator is tuned to, with the mode selector, polyphony and latch stacked in its right-hand column. The INPUT block in the middle is everything that happens to the spectrum before it reaches the resonator, with the STFT analysis settings stacked in its right-hand column. The right-hand block is split in two: its top row is TUNING, which reshapes the harmonic series and can snap it to a scale and its lower two rows are the RESONATOR itself, the mask, the decay and the damping, together with the Freeze switch and, in the instrument build, the PLAY SYNTH and POLY switches.

6.36.1 VOICING

The voicing block sets which pitches the resonator is tuned to. Three voicing modes are selectable from the Mode dropdown at the top of the block's right-hand column: FIXED, MIDI and TRACK. The choice determines which of the other voicing controls are active and how the V1 NOTE knob's display behaves.

FIXED is the default. The four V1-V4 knobs drive the voicing directly. V1 sets an absolute pitch and V2 / V3 / V4 are semitone offsets above or below it.

MIDI replaces the four manual voices with notes received from the host. V1-V4 are greyed out, GLIDE / PB / Polyphony / LATCH light up and become active. In MIDI mode you get up to twelve simultaneous polyphonic voices with FIFO voice stealing when the polyphony cap is exceeded. In the effect variant of the plugin, MIDI events are only routed by hosts that support it; in Logic Pro this means loading the plugin as a MIDI-controlled effect with the source synth side-chained in and both tracks record-ready.

TRACK listens to the incoming audio and finds its pitch and feeds that pitch into the voicing instead of a fixed reference. The detection is done in the time domain, by finding the interval at which the incoming waveform repeats, which is far better at the bottom of the range than reading a spectrum is: a bass note at 40 Hz is found within a fraction of a cent and a bass whose fundamental is missing altogether, as many are, is still found at the pitch you hear. Frames where the input is not clearly periodic, a transient or noise, do not move the pitch and a single-frame error never gets through. V2, V3 and V4 still apply as semitone offsets from the detected pitch, so the engine produces a moving harmonic field that follows whatever pitch the input is at. V1 in TRACK mode becomes a signed semitone offset around the detected pitch rather than an absolute note: V1 at 60 means zero offset and the resonator tracks the detected pitch exactly, V1 at 72 transposes it up an octave, V1 at 48 down an octave. The knob's display reflects this, see V1 NOTE below. TRACK follows one pitch. For a source with several notes at once there is a second tracker that follows every pitch it hears, used only when the field is playing the synthesiser; see the POLY switch and section 6.36.7.

V1 NOTE - The fundamental MIDI note number of Voice 1, from 0 to 127. Default is 60. The displayed text adapts to the current voicing mode. In FIXED mode the value reads as a note name following the Yamaha/Roland convention where MIDI 60 is C3, MIDI 72 is C4, MIDI 48 is C2 and so on - the formula is octave = (MIDI / 12) - 2, so middle C is shown as C3 not C4. In TRACK mode the value reads as a signed semitone offset from 60: "0 st" at the centre, "+12 st" twelve semitones above and "-12 st" twelve semitones below. In MIDI mode the knob is greyed out and the value reads as a note name for reference. Voice 1 is always active in FIXED and TRACK modes and serves as the reference pitch for the V2 / V3 / V4 offset knobs.

V2 OFF / V3 OFF / V4 OFF - Per-voice pitch offsets in semitones relative to V1, from -36 to +36 semitones. Each of V2 / V3 / V4 has a small enable LED in the top-right corner of its knob bezel - click the LED to add or remove that voice from the active set. Defaults are V2 at -12, V3 at 0 and V4 at +12. With V2-V4 all enabled you get a four-voice harmonic field; with only V1 active you get a single-voice resonator. The knobs snap to whole semitones.

GLIDE - Portamento time between MIDI notes, in milliseconds (0 to 2000 ms). Only active in MIDI mode. At 0 ms new notes snap instantly; above zero a one-pole interpolator smooths the transition, with the value being the approximate 99 percent settling time so 100 ms produces a glide that completes in about a tenth of a second. The glide is per-voice - chord voicings glide independently to their new pitches rather than as a parallel block.

PB - Pitch bend range in semitones (0 to 24), interpreted as the maximum reach when the host pitch wheel is at full deflection. Default is 2. Only active in MIDI mode. The bend is applied to all active MIDI voices simultaneously.

FINE - Fine pitch trim applied to the resonator's harmonic series as a whole, in cents (-100 to +100). Useful when you need to detune the entire field by a few cents to sit better with other tracks or with the song's tuning. Distinct from V2-V4 OFF which sets pitch offsets between voices - FINE shifts every voice's harmonics together by the same amount.

UNISON - Number of sub-voices per active voice, from 1 to 7. Default is 1. Higher values produce a thicker more chorused sound by overlaying multiple slightly-detuned copies of each voice's harmonics. CPU cost scales linearly with unison count.

DETUNE - Detune spread between unison sub-voices, in percent (0 to 100). At 0 percent all unison voices sit on top of each other and UNISON has no audible effect; increasing DETUNE spreads the sub-voices across a band centred on each voice's fundamental, producing the classic chorus / supersaw thickening. The maximum spread at 100 percent is about plus or minus 25 cents.

The right-hand column of the voicing block holds the mode dropdown and two MIDI-only controls stacked vertically:

Mode dropdown - **FIXED / MIDI / TRACK** The default is FIXED. In the effect variant of the plugin the MIDI entry in the dropdown is greyed out, since the effect-build cannot reliably receive MIDI unless the host wraps it as a MIDI-controlled effect.

Polyphony combo - Maximum simultaneous MIDI voices, selectable between 4, 6, 8 (default), 10 and 12. When more notes are played than the cap allows, the oldest active note is stolen to make room for the newest (FIFO voice stealing). Active only in MIDI mode.

LATCH (toggle) - When on, MIDI note-off events are ignored - notes stay active until they are stolen by a new note or until the transport stops. Combined with FIFO stealing this gives chord-stacking behaviour where every new note pushes out the oldest. Useful for building up a slowly-evolving drone from a sequence of played notes without having to keep them held down. Active only in MIDI mode.

6.36.2 INPUT PIPELINE

The input block is the chain of spectral processors that operate on the input's frequency-domain representation before it reaches the resonator. Each stage runs once per STFT frame and is independent of the others; setting them all to neutral values is equivalent to bypassing the block entirely. None of them affect what the pitch trackers hear: the trackers read the input as it arrives, before this block.

The block's first column holds the Filter Type selector as a vertical stack of three buttons, with GATE and SHIFT beneath it. The next two columns hold the remaining knobs and the right-hand column holds the STFT analysis settings, Block and Hop.

Filter Type - LP / HP / BP - Three buttons stacked vertically in column one. Click a button to engage that filter type and silence the others; click an active button a second time to turn the filter off entirely (no button lit = filter bypassed). This is the only segmented group in the engine that allows an all-off state - the other segmented groups (Harmonic Filter, Stretch Mode) always have exactly one option selected.

FREQ - Filter centre frequency in Hz (20 to 20000) for BP mode or the cutoff frequency in Hz for LP and HP modes. Default is 1000 Hz. Frequency-adaptive display (shows as Hz below 1 kHz, kHz at and above).

BW - Filter bandwidth in semitones (1 to 120) for BP mode or transition-band width for LP and HP modes. Default is 24 semitones (two octaves). A BP filter with BW set to 120 semitones (ten octaves) is effectively wide-open and is treated as a filter bypass internally to save CPU. For LP and HP modes BW controls how steep the cutoff is - small values give a sharp brick-wall response and large values produce a gentle gradual rolloff.

SHAPE - Continuous control over the filter's edge shape, from 0.0 (rectangular brick-wall) through 1.0 (raised-cosine smoothing) to 2.0 (Gaussian smoothing). Default is 1.0. Higher values produce gentler, more natural-sounding edges with less ringing; lower values produce sharper more surgical filtering.

GATE - Spectral gate threshold at 1 kHz reference, in dB (-90 to 0). At -90 dB (the default) the gate is effectively off and every spectral bin passes through. Above -90 dB the gate silences any bin whose magnitude falls below a per-bin threshold curve - the curve follows a -3 dB/oct slope so low frequencies need to be louder to pass than high frequencies, matching the rough shape of pink noise spectral density. Use this to keep quiet noise out of the resonator on sustained material or as a creative "spectral mute" that follows the input's dynamics.

SHIFT - Spectral frequency shift in semitones-equivalent (-24 to +24), interpreted at the 1 kHz labelling reference. Default is 0. A positive shift translates every spectral bin upward by the same number of FFT bins (computed as the round-trip-frequency equivalent at the reference frequency); a negative shift translates downward. Unlike a pitch shifter this is a *frequency* shift rather than a *pitch* shift - the harmonic relationships are not preserved and the output becomes inharmonic at non-zero values. Subtle settings (± 1 to ± 3) produce metallic detuning; larger settings produce striking inharmonic effects.

BLUR - Magnitude blur kernel width in FFT bins (0 to 32). At 0 (the default) no blur is applied. Above zero, each bin's magnitude is replaced by a weighted average across a window of $(2 \times \text{BLUR} + 1)$ neighbouring bins. The effect is a spectral smearing - sharp peaks become softer rounded humps and the resonator excites a wider region around each harmonic. Useful for taking the edge off material that is too "pure" and for producing washier softer-sounding resonator output.

BLR SH - Blur kernel shape, same 0.0 to 2.0 scale as the filter SHAPE (rectangular / raised-cosine / Gaussian). Default is 1.0. Affects the character of the blur - rectangular gives a flat box average, Gaussian a smooth bell-shaped weighting.

DENSITY - Probability that each harmonic of each voice is excited on any given frame, in percent (0 to 100%). Default is 100% (every harmonic excited every frame). Below 100% the engine rolls a dice for each harmonic each frame and skips excitation if the roll fails - the result is a sparse "twinkling" texture where individual harmonics flash in and out at random rather than ringing continuously. At very low values (10-30%) you get a sparkling shimmer that decays into the resonator's tail; at moderate values (50-70%) you get a subtle spectral animation; at 100% you get the dense classical resonator sound.

The rightmost MISC column on row 2 holds the STFT analysis settings. These two combos affect the engine's CPU cost, latency and tonal character - *changing them mid-session causes a brief audio gap while the engine reallocates its internal buffers and the host re-reads the new latency value.* Save your work before experimenting.

Block Size combo - STFT analysis block size in samples (256, 512, 1024, 2048, 4096). Default is 1024. Larger sizes give better frequency resolution and a smoother more pitched sound but at the cost of higher latency and slower response to transients. 1024 is a good general default; 2048 or 4096 for pad-like material where transient response doesn't matter; 256 or 512 for percussive material where you want the resonator to track the attacks.

Hop combo - STFT hop factor (2, 4, 8, 16). Hop factor is the overlap ratio N/H where N is the block size and H is the hop length - a hop of 4 means each block overlaps the previous by 75% (one-quarter of a block advanced per analysis). Default is 4. Higher values give smoother output at the cost of more CPU; lower values are cheaper but can introduce audible buzzing on tonal material. 4 is recommended for most uses; 8 for the cleanest output; 2 only if you're CPU-bound.

6.36.3 TUNING AND RESONATOR

The right-hand block is the heart of the engine. Its top row, TUNING, decides where the harmonics of each voice fall: how many there are, which of them are used, how the series is reshaped and whether it is snapped to a scale. Its lower two rows, RESONATOR, decide how the resonator responds at those positions: the width and shape of each harmonic's excitation band and how long each holds its energy.

Reading the tuning row left to right: the Harmonic Filter selector, the HARM count knob, the LAW selector, the STRETCH knob and the ROOT and SCALE combos of the quantiser stacked in the last column. Reading the resonator rows: MASK BW, MSK SH, DECAY, HF DMP and LF DMP across the middle row, with FRZ and, in the instrument build, PLAY SYNTH and POLY along the bottom.

Harmonic Filter - ALL / EVEN / ODD - Three buttons stacked vertically in column one. Default is ALL (every harmonic excited). EVEN excites only the even-numbered harmonics (2nd, 4th, 6th, ...) - this produces a hollow "clarinet-like" or "square-wave-like" timbre where the fundamental is suppressed and the spectrum is dominated by overtones of the same family. ODD excites only the odd-numbered harmonics (1st, 3rd, 5th, ...) - this produces a brighter "sawtooth-like" or "string-like" timbre. The active mode is always one of the three (unlike Filter Type this is a strict radio group with no off state).

HARM - Number of harmonics excited per voice (1 to 64). Default is 16. A single voice with HARM at 64 spans roughly six octaves above its fundamental (assuming Linear stretch at 100%); restricting HARM to a smaller number gives a narrower more focused spectral footprint. Harmonics above Nyquist are automatically silenced.

LAW - LIN / FLET / OCT - Three buttons stacked vertically, choosing how the STRETCH knob reshapes the harmonic series. Default is LIN. Under LIN the harmonics sit at $f_0 \times (1 + (n-1) \times \text{stretch})$, so STRETCH is the interval between neighbouring harmonics: at 100% this is the natural series, at 200% every other harmonic, at 50% an extra one between each pair. Under FLET the harmonics follow the stiff-string formula $f_0 \times n \times \sqrt{1 + B \times n^2}$ with B derived from STRETCH, which is the formula used to model piano strings and it makes the upper harmonics progressively sharp as a real string's are. Under OCT the series is bent, in the log domain, between the natural series and a series of octaves: at 100% it is the natural series, at 200% every harmonic is exactly an octave above the last and between the two the series climbs faster than the natural one. Below 100% it is squashed the other way, climbing ever more slowly. OCT is the law that turns a harmonic ring into something organ-like or bell-like. Under every law the fundamental itself stays where the voicing put it.

STRETCH - The amount by which the chosen law reshapes the series, in percent (0 to 300). Default is 100, which is the natural series under every law. Under LIN, 50% squashes the harmonics together (the 2nd at 1.5f instead of 2f) and 200% spreads them apart (the 2nd at 3f, the 3rd at 5f). Under FLET the value scales the stiff-string inharmonicity: at 100% the upper harmonics are noticeably sharp like a real piano string, at 200% more so, like a bell. Under OCT, 200% is pure octaves and values below 100 squash the series so its upper members crowd together.

ROOT and SCALE - The scale quantiser, two combos stacked in the last column of the tuning row. With SCALE set to anything other than Off, every harmonic above the fundamental is snapped to the nearest note of that scale on that root, after the LAW and STRETCH have placed it. The fundamental itself is never moved. Twelve roots are offered and eleven scales: Chromatic, Major, Minor, Dorian, Phrygian, Lydian, Mixolydian, the major and minor pentatonics, Blues and Whole tone. Because the snap happens after the stretch, the tuning controls decide which notes land in the chord and the scale decides which chord it is: set STRETCH somewhere unusual, choose a scale and the field stops ringing at a series and starts ringing at a chord in your key. With SCALE at Off the series is exactly as the law places it.

MASK BW - Width of each harmonic's excitation band in FFT bins (1 to 32). Default is 4. At small values each harmonic is a narrow needle on the spectrum and the resonator output sounds pure and ringing; at large values each harmonic is a broad band and the resonator output sounds washier and more diffuse. The effective musical width depends on the STFT block size - at Block

1024 / 48 kHz, one bin is roughly 47 Hz, so MASK BW of 4 corresponds to about ± 94 Hz around each harmonic.

MSK SH - Mask shape, same 0.0 to 2.0 scale as Filter SHAPE and BLR SH (rectangular / raised-cosine / Gaussian). Default is 1.0. Gaussian produces the smoothest most natural-sounding harmonic excitation; rectangular gives a sharper more synthetic character.

DECAY - The -60 dB decay time of each resonator bin, in milliseconds (10 to 60000). This is the headline control of the engine. At 10 ms (the minimum) the resonator decays almost as fast as the input - you hear it as a subtle spectral colouration. At 1000 ms (the default) each bin holds its energy for about a second. At 10000 ms the resonator rings for a full ten seconds. At 60000 ms (the maximum) the resonator is effectively infinite - combine with FREEZE for true endless sustain. Modulating DECAY produces dramatic textural changes - try a slow LFO between 200 ms and 5000 ms for breathing pads.

HF DMP - High-frequency decay damping frequency in Hz (100 to 20000). Default is 10000. Bins above this frequency decay faster than bins below it, producing the natural "darker tails" you hear in real resonators and reverbs. Lower values give darker shorter top-end; higher values preserve more sparkle in the tail.

LF DMP - Low-frequency decay damping frequency in Hz (20 to 2000). Default is 100. Bins below this frequency decay faster than bins above it, removing low-end buildup from the resonator tail. Especially useful when you have multiple voices stacked and the low end accumulates into muddy rumble; bringing LF DMP up to around 200 to 400 Hz cleans this up while keeping the body of the sound intact.

Along the bottom row of the resonator block:

FRZ (toggle) - Spectrum Freeze. When engaged, the engine stops reading new input and the current spectral state is held indefinitely. Combined with a long DECAY this produces infinite sustain of whatever was in the resonator at the moment Freeze was activated. Useful for capturing a single chord or texture and holding it as a drone while you continue to play. Disengage to resume reading input.

PLAY SYNTH (toggle, instrument build only) - Sends the pitches the field is tuned to as notes to the synthesiser in the topmost layer, so that the field, listening to a side-chain or to another layer, plays the synthesiser. Off by default. It is described in full in section 6.36.7.

POLY (toggle, instrument build only) - With PLAY SYNTH on in TRACK mode, chooses which of the two pitch trackers drives the synthesiser. Off follows one pitch, which is right for a bass line or a lead. On follows every pitch the field can hear, which is right for a pad or a chord. Off by default. See section 6.36.7.

6.36.4 THE VISUALISER

The visualiser sits below the control panel and refreshes at 30 frames per second. It is an educational display, built to help you understand what the engine is doing rather than to be a precision measurement tool. The display is a single 2D panel with a logarithmic frequency axis from 20 Hz to 20 kHz running left to right and a magnitude axis running bottom to top, scaled to comfortably accommodate the wide dynamic range of raw FFT bin magnitudes plus the gain build-up of the SMD resonator without clipping at the top edge. Grid lines mark octave boundaries at the standard musical frequencies and frequency labels appear at 100 Hz, 1 kHz and 10 kHz across the bottom.

Several layers are drawn back to front:

1. Filter passband as a shaded coloured region - the current shape of the Filter Type / FREQ / BW / SHAPE settings, drawn so you can see at a glance how much of the input is being preserved or attenuated by the filter stage. Disappears entirely when no LP/HP/BP button is lit (filter bypassed).

2. Past-spectrum wireframes in a faint version of the engine's accent colour - thin skyline traces of recent SMD output spectra, drawn lifting upward and slightly compressed horizontally to give a receding-perspective impression of spectral history. The wireframes fade with age so the oldest frames are nearly invisible. Useful for seeing how the resonator tail evolves over time and for spotting where energy is accumulating or shedding.

3. Newest-slice filled spectrum - the current SMD output as a filled coloured region anchored at the bottom of the display, with a horizontal rainbow gradient running from deep blue at 20 Hz through green and yellow to pink-red at 20 kHz. This is the foreground layer and what you watch when adjusting DECAF, HF DMP and LF DMP. The fill is partially transparent so the past wireframes show through behind it.

4. Harmonic mask positions as small vertical tick marks along the top edge - one mark per active voice's harmonic, accounting for the ALL/EVEN/ODD selection, all three laws and the scale quantiser and drawn from exactly the same placement the resonator uses, so the picture cannot disagree with the sound.

5. Voice fundamental markers as small dots along the bottom axis at each active voice's f-zero.

A FREEZE badge in the top-right corner of the display indicates when the Freeze toggle is engaged.

The visualiser reads from a published-state buffer that the audio thread updates once per STFT hop. It does not run its own FFT - what you see is exactly what the audio thread computed, with the only difference being that the visualiser samples that buffer at 30 Hz so very fast spectral changes between updates are not visible.

6.36.5 FACTORY SUB-PRESETS

Eight factory sub-presets are provided as starting points, available from the sub-preset dropdown at the top of the panel. Each is a snapshot of every parameter of the engine captured at a setting that demonstrates a particular character. The sub-presets predate the quantiser and the PLAY SYNTH switches, which they leave at Off.

Resonant Choir - Four voices stacked at the root, major 3rd, 5th and octave; a broad band-pass at 2 kHz with smooth Gaussian edges; 32 harmonics per voice; 3-second decay. The default starting point for choir-like washes.

Tape Halo - Single voice at the root; narrow band-pass at 800 Hz; 20 harmonics in Fletcher stretch mode for warm inharmonic colouration; light spectral blur for tape softness; 800 ms decay; warm high-frequency rolloff. Suited to vintage tape-style spectral wash effects.

Crystal Hall - Three voices stacked (root, 5th, octave); filter off (wide-open spectrum); 48 dense harmonics; 20-second decay with Freeze engaged for indefinite sustain; high-frequency damping relaxed to preserve sparkle. The "infinite hall" preset.

Pluck Tail - Single voice; low-pass filter to keep things warm; 12 even-only harmonics for a hollow plucked-string character; 400 ms decay for a quick natural fade. Use on percussive material to add a brief pitched bloom after each hit.

Atmosphere - Four voices stacked across two octaves; broad band-pass at 1.5 kHz; the full 64 harmonics per voice; heavy spectral blur; low density (50%) for sparkling motion; 8-second decay; strong low-frequency damping to keep the field clean. A slow wide atmospheric texture.

Spectral Drone - Single low drone voice; 300% stretch for stretched inharmonic spectra; narrow mask for purity; 15-second decay with Freeze engaged. Single-pitch drone with endless evolving texture.

Ghost Vocal - Three unison voices at the root with 15% detune spread for chorus motion; odd-only harmonics for a vocal-formant character; Gaussian mask shape for smoothness; 2-second decay. Useful for ghosted vocal-like resonator output.

Inharmonic Shimmer - Three voices stacked; 40 partials per voice; Fletcher stretch at 180% for pronounced inharmonicity; 600 ms decay so the texture stays present rather than washing into a long tail. Bright sparkling motion suited to ambient and cinematic material.

6.36.6 WORKING WITH RESONANCE FIELD

The character of the engine depends as much on what you feed it as on how you set it. A few starting points:

Drum loops: set GATE to about -50 dB so quiet hat noise doesn't trigger the resonator, then sweep DECAY from short (300 ms) to long (3000 ms) - the kick and snare become pitched accents and the loop turns into a pitched melodic pulse. Try the Resonant Choir or Crystal Hall preset as a starting point and adjust V1 NOTE to fit the song key.

Sustained pads or strings: drop DENSITY to 50-60% to introduce spectral sparkle, raise MASK BW to 6 or 8 for a washier texture and let DECAY sit between 1500 and 4000 ms. Atmosphere is a good starting point for this style of use.

Vocals or speech: drop UNISON to 1 and set V2/V3/V4 OFF to widely spread intervals (e.g., +7, +12, +19) to produce a harmonised effect that follows the singer's articulation. Ghost Vocal is the natural starting point - try also Tape Halo on whispered or breathy material.

Sound design / cinematic stings: load Inharmonic Shimmer, set DECAY to a long value (5000 ms or more), engage FREEZE and play a chord on the host MIDI controller (after switching to MIDI mode). The Freeze captures the chord and holds it as a sustained inharmonic bell that you can then process further down the chain - modulate STRETCH slowly via an LFO for a continuously evolving inharmonic drone.

As a MIDI-controlled effect in Logic Pro: load the Effect variant into the Instrument slot of a Software Instrument track. Place the source synth on a separate track and side-chain its output into the NeuralDream track. Mute the source synth's main output (so only Resonance Field's processed output is heard) and arm both tracks record-ready. With the engine in MIDI voicing mode, played notes drive the voicing while the side-chained synth audio excites the resonator. This is the workflow that gives the most expressive control over the engine in a live or composed setting.

Latency note: the engine's latency is dominated by the STFT block size and hop factor. At the default 1024-sample block with hop factor 4 (256-sample hop) the latency is approximately 1280 samples - about 27 ms at 48 kHz. This is reported to the host so DAW plugin delay compensation handles timing alignment automatically. If you reduce the block size to 256 the latency drops to about 8 ms; increasing to 4096 pushes it to about 100 ms. Most musical uses tolerate the default without issue; for tight rhythmic playing where the resonator needs to track transients precisely, drop the block size at the cost of some spectral resolution.

6.36.7 PLAYING THE SYNTH FROM RESONANCE FIELD

In the instrument build, Resonance Field can play the synthesiser. The field listens to a sound, finds the pitches in it and sends them to the synthesiser in the topmost layer as notes, exactly as if they had been played on a keyboard. The synthesiser then plays them through whatever oscillators, filters and envelopes the patch has. Play a bass line into the side-chain and the synthesiser follows it; feed it a pad and it plays the chords. The field's own resonated sound is still produced on its layer, so you can use both together or turn the field's layer down and keep only the synthesiser.

Setting it up. Layer 1 holds the synthesiser as usual. Add a second layer and set its source to the audio you want the field to listen to: SideChain if it is coming from another track in your DAW or Layer Output if it is another layer inside the plugin. Add Resonance Field to that layer and set its Mode to TRACK. In Logic, choose the track or bus you want to send in the Side Chain menu at the top right of the plugin window as well; setting the layer's source tells NeuralDream where to look and the Side Chain menu tells Logic what to send and both are needed. Then press PLAY SYNTH. Every note is sent at velocity 100.

MONO or POLY. The POLY switch beside PLAY SYNTH chooses which of two trackers drives the synthesiser and the choice matters more than any other setting. With POLY off the field follows one pitch: the time-domain tracker finds the interval at which the incoming waveform repeats and that is the note. It is right for a bass line, a lead, a vocal or any single line and it is confident enough that a change of note is taken only when it is sure, so a chord passing over a bass does not make it flutter. With POLY on the field follows every pitch it can hear: a note is a pitch that appears in the spectrum and persists and it ends when that pitch has gone. A chord gives a chord, up to six notes at once. It is right for a pad, a chord part or anything with more than one voice.

Why it is a switch and not automatic. A bass often has no energy at its fundamental at all; what the ear hears is the common period of its harmonics. Seen from the spectrum that is indistinguishable from a chord voiced in octaves and fifths, which really is a harmonic series. Any rule that reads one correctly reads the other wrongly, so the choice is yours. Off for one line, on for several.

Settings that matter. Set Block to 4096 for POLY: it is the window that lets the tracker tell adjacent notes apart, at the cost of a fraction of a second more latency than 2048. Switch V2, V3 and V4 off unless you want the synthesiser to play harmonies: with POLY off, every enabled voice is sent as a note, so a voice at minus twelve semitones puts an octave under everything the field hears. Make sure the synthesiser itself is polyphonic. Presets saved before POLY existed load with it off; press it and save the preset again.

How notes start and end, in POLY. A pitch becomes a note once it has been present for a few frames, so a transient never triggers one and a pad note swelling in over a tenth of a second is caught once it has settled. A note is held for as long as its pitch is present: it ends when it has fallen to thirty decibels below its own peak or when it has become insignificant next to the notes that have replaced it, so a note decaying on its own is one note from attack to silence and a chord change releases the old chord promptly. A note that has ended cannot come back on its release tail; only a genuine re-strike, arriving nearly as loud as the note was, retriggers it and a re-strike of a note that is still held retriggers it too. Beating between detuned oscillators and a note bending down a semitone as it releases, are recognised for what they are and do not start new notes.

What it cannot do. Octaves and fifths collapse: a C3 with a C4 and a G4 above it is, acoustically, a C3 with its second and third harmonics and the tracker will play the C3, because it cannot tell a voicing from a timbre. Voicings in thirds, sixths and sevenths come back as chords. Notes a semitone apart need to be high enough for the window to separate them: at Block 4096 that is from about C5 upwards, a whole tone from about C4 and wider intervals below. And the source must be tonal: a clean synth, a piano, a plucked or bowed string work well; a distorted guitar chord, a heavily vibrato vocal or anything noisy will not.

6.37 FLOW

Flow is a spectral freeze. It captures the momentary spectrum of the incoming audio and rebuilds it continuously with fresh phase, so a passing sound becomes an endless, evolving pad that holds for as long as you like. It works on a single audio source and, like the other STFT engines, reports its analysis latency to the host so plugin delay compensation keeps everything aligned.

Flow has two modes, chosen by the MELODIC toggle:

PHASING (MELODIC off) - the default. Flow freezes the magnitude spectrum and randomises the phase, rebuilding the sound grain by grain with new phase each time. The result is airy and shimmering with no fixed pitch - ideal for atmospheres, drones and ambient wash.

MELODIC (MELODIC on) - Flow tracks the strongest partials as real sinusoids and randomises only the noise residual around them. This holds a stable, pitched chord rather than an airy cloud, so a frozen note or chord keeps its pitch and harmony.

The remaining controls:

FREEZE - latches the spectrum. With FREEZE on, Flow stops following the input and holds the captured spectrum indefinitely, until you release it. FREEZE is a modulation target, so you can gate the freeze from an envelope, an LFO or a MIDI control.

DECAY - how long a captured spectrum is held when FREEZE is off, from 0 to 100%. At 0 the freeze follows your playing closely, refreshing as the input changes; at 100 it holds for roughly thirty seconds before fading, for long sustained pads.

WIDTH - stereo decorrelation of the rebuilt sound, from 0 to 100%. At 0 the two channels are identical (mono); as WIDTH rises the left and right rebuilds become independent, opening the freeze into a wide stereo field.

MIX - the dry/wet blend, from 0 to 100%. The dry signal is delay-matched to the wet so the two stay phase-aligned when blended.

Flow prefers the full visualiser, where the frozen spectrum is drawn live. Its accent colour is violet.



6.38 COMPRESSOR

A stereo dynamics compressor built on the feed-forward, log-domain design set out in Giannoulis, Massberg and Reiss, "Digital Dynamic Range Compressor Design - A Tutorial and Analysis" (Journal of the Audio Engineering Society, 2012) - the modern reference for smooth, musical, artefact-free compression. The level detector runs in the decibel domain with decoupled attack and release smoothing, which lets fast attacks grab transients without the pumping and crackle of linear-domain designs.

THRESH - the level, in dB, above which gain reduction begins.

RATIO - how much of every decibel above the threshold is allowed through, from 1:1 to 40:1. Gentle levelling sits around 2:1 and firm control around 8:1. The highest settings behave as a limiter and are there for holding a sound at a steady level: a part whose level swings by 24 dB as you move the mod wheel comes out varying by about 1.2 dB at 20:1 and about 0.6 dB at 40:1.

KNEE - the width, in dB, of the soft-knee region centred on the threshold. At 0 the transfer curve has a hard corner; wider values blend gain reduction in gradually across the knee, using the quadratic interpolation from the reference design.

ATTACK / RELEASE - the detector time constants, with the live values always readable at the top of the visualiser.

GAIN - make-up gain after compression.

The visualiser presents the compressor's law, its present moment and its recent history on one panel. The main pane draws the static transfer curve - input level along the horizontal, output up the vertical, both axes captioned and gridded in 12 dB steps - with the threshold marked and the current ATTACK and RELEASE times printed above the curve in milliseconds. A live operating point rides the curve at the current input level: watch it climb the unity slope, bend through the knee and travel the compressed region as the music plays. Note that the curve's horizontal axis is level, not time; the time constants show themselves in the operating point's motion - how fast it dives onto the compressed slope, how it relaxes back.

To the right, three calibrated meters show IN, GR and OUT. Input and output side by side are the before-and-after; the gain-reduction meter between them fills top-down by the dB currently being removed, on a 24 dB scale ticked every 6 dB with its own figures, while the shared IN/OUT scale is numbered every 12 dB down to -60. All three use hardware ballistics - instant rise, smoothed fall - and the GR meter carries a peak-hold tick that lingers for a moment at the deepest recent reduction.

Beneath the curve runs the GR/TIME strip: gain reduction traced against real seconds, a four-second window scrolling with the newest moment at the right, marked every second and referenced at -6, -12 and -18 dB on the same top-down scale as the GR meter. The engine feeds it at a fixed 4-millisecond resolution, capturing the deepest reduction within each interval so a fast transient's clamp can never slip between display samples - and because the interval is fixed in real time, the axis reads true in any host at any buffer size. This is where the time constants become visible as shapes: an attack draws a cliff whose steepness you set with ATTACK, the recovery draws a slope governed by RELEASE and pumping at musical rates draws as literal waves. The engine's panel artwork sets the colour family throughout: lattice green for input, arc yellow for gain reduction and the time trace, beige for output.

6.39 DARKHOLE

A vast, gravitational reverb built from a long chain of modulated allpass filters — sixteen in series, then eight more per channel — inspired by the celebrated long allpass-chain reverbs of Eventide's studio processors. Where the FDN and Lustre reverbs model rooms and plates, DARKHOLE models nothing in particular: it is a space that could not exist, capable of gentle three-second halls at one setting and infinite, slowly rotating cathedrals of sound at another. It is the engine to reach for when the reverb is meant to be heard as an instrument rather than as an acoustic.

DARKHOLE is the only NeuralDream reverb whose size and pitch character can be pushed into deliberately impossible territory and it rewards extremes: very long pre-delays, size settings beyond the natural range and feedback held just below runaway.

GRAVITY – the allpass coefficient, bipolar from -100 to +100. This is the pull of the field: it sets how strongly each stage smears its input and therefore how dense and how diffuse the space becomes. Negative and positive values shape the smear differently, so it is worth sweeping the whole range rather than treating the centre as neutral.

SIZE – the length of the delay lines, from -10 to 120. Low values give tight, metallic spaces; the upper reaches open into something enormous and slow. Changes are smoothed, so sweeping SIZE while sound is passing produces a long gravitational glide rather than a click.

PRE-DLY – pre-delay before the reverb begins, from 0 to 2 seconds. Generous by design: at long settings the dry sound is fully finished before the space answers it.

FEEDBACK – how much of the output is returned to the input, from 0 to 110%. Below 100% the tail decays; the last stretch of the range is the near-infinite zone, where the sound sustains almost indefinitely. It cannot be driven into runaway – the feedback path is protected – but it can be held on the edge of forever.

MOD and **MOD RT** – depth and speed of the modulation running inside the allpass chain. A little movement prevents the metallic ringing that a static allpass network is prone to; a lot produces obvious pitch drift in the tail, which suits pads and drones.

FREEZE – infinite hold. The input is muted, modulation stops and whatever is in the space stays there, ringing on until you release it. Combine with a long PRE-DLY to capture a single chord and hold it while you play over the top.

REV – reverse mode, in which the tail swells toward the transient rather than decaying away from it.

6.39.1 THE LO CUT AND HI CUT CONTROLS

These two controls are filters inside the feedback path and they behave differently from the shelving EQ found on the other reverbs. They can only ever remove energy from the tail – never add it – so no combination of settings can push the reverb into runaway. This is deliberate: a boost inside a feedback loop multiplies with every pass.

Their resting positions are at opposite ends of their travel, because each one closes in from its own end of the spectrum:

LO CUT is a high-pass corner, from 20 Hz to 1 kHz. Fully anticlockwise (20 Hz) is OFF – nothing is removed. Turning clockwise raises the corner, progressively stripping low end out of the tail. Use it to stop a large space from clouding the bottom of a mix.

HI CUT is a low-pass corner, from 20 kHz down to 500 Hz. Fully clockwise (20 kHz) is OFF – nothing is removed. Turning anticlockwise lowers the corner, progressively darkening the tail. Because the filter sits in the feedback path, each successive pass is darker than the last, which is how natural rooms behave.

RESO – resonance at both filter corners described below, from a gentle, critically damped slope at 0% to a pronounced emphasis at the corner frequencies at 100%.

So at their default positions – LO CUT hard left, HI CUT hard right – both filters are bypassed and DARKHOLE sounds exactly as it does with no contouring at all. Moving either knob inwards begins to narrow the band that survives inside the loop and RESO sharpens whichever corners are in play. The controls are for shaping the sound inside the space; they are not a tone control on the output and they will not change the fundamental character of the reverb.

6.39.2 THE VISUALISER

A gravity well. Stars drift through deep space and bend toward the centre as they approach it, reddening as they fall; an accretion disc of photon streaks orbits the event horizon, the inner orbits sweeping faster than the outer ones and burning from violet at the rim through magenta and amber to white at the innermost edge. A ring of bent light marks the horizon itself, inside which there is nothing at all.

The display is driven by the engine: the disc spins faster as the output level rises, grows deeper and wider with FEEDBACK and GRAVITY, wobbles with MOD and burns hotter with RESO. Engaging FREEZE stops the field dead and marks it as frozen.

6.40 FILAMENT

Filament re-tunes a sound's harmonics without moving its note.

That sentence is worth reading twice, because it describes something neither of the familiar tools can do. A pitch shifter multiplies every partial by the same factor and takes the fundamental with it - the note changes. A frequency shifter adds a constant to every partial and takes the fundamental with it too - the note changes and the harmonic relationships collapse. Filament is the third member of that family: the fundamental is a fixed point and everything above it slides. The note you played stays exactly where you put it, while the material it appears to be made from changes underneath it.

The ear reads this as a change of substance rather than a change of pitch. Pull the harmonics together and a bright string tightens into something reedy and hollow. Push them apart and the same note acquires the shimmer of struck metal - a bell, a tine, a plate. Nothing has been detuned and nothing needs to be played differently; the note is the same note, made of a different metal.

6.40.1 HOW IT WORKS

Every harmonic sound is a ladder: a fundamental with rungs above it at whole-number multiples. Filament finds that ladder, *respaces* the rungs according to your settings and rebuilds the sound around them - leaving the bottom rung exactly where it was.

Because the effect is defined relative to the fundamental, it has to know which note is sounding. It finds this for itself by default and follows you as you play.



6.40.2 THE CONTROLS

STRETCH (-100 to +100%, default 0) - the heart of the engine. It spreads the harmonics apart or pulls them together and the note itself never moves. Positive values stretch the ladder, which is the behaviour of real struck and stiff materials: bells, tines and piano strings all have partials that sit progressively sharp of the ideal series. Negative values compress it, producing hollow, reedy, faintly detuned timbres that have no natural equivalent at all. At zero the engine is transparent.

CURVE (0 to 100%, default 0) - blends the *shape* of the spread. At 0 the respacing is even, which is the character of bells and gongs. At 100 it follows the stiff-string law, which is the character of tines and struck piano wire. The difference is subtle on small STRETCH values and pronounced on large ones; it is the control that decides whether a stretched sound reads as *bell* or as *electric piano*.

ORDER (-100 to +100%, default 0) - thins one half of the series. Positive values reduce the even harmonics, which hollows the sound toward the woody, clarinet-like character of an odd-harmonic spectrum. Negative values thin the odd ones instead. It is a strong control: small amounts colour, large amounts transform.

FOCUS (0 to 100%, default 100) - how much of the non-tonal content stays where it is. Real instruments are not only harmonics: there is breath, bow noise, key clatter and room. At 100 Filament re-tunes only the tonal ladder and leaves everything else untouched, which keeps the sound believable. Lower values let the noise be dragged along with the harmonics, which is less natural and often more interesting.

ROOT (27.5 Hz to 1200 Hz, default 110 Hz) - the note the harmonics are measured from, used when AUTO is off.

AUTO (on/off, default on) - find the note automatically. Leave it on for single lines and monophonic material. Turn it off and set ROOT by hand, for chords and dense material, where there is no single fundamental to find and a fixed reference gives a far steadier result.

GLIDE (0 to 2000 ms, default 60 ms) - how quickly the engine follows a new note. Short values are accurate and track fast playing cleanly. Long values smear one note's harmonic frame across the next, which is a usable effect in its own right on slow, sustained material.

MIX (0 to 100%, default 100%) - dry/wet blend. The dry path is delay-matched to the wet, so blending does not produce phase cancellation or comb filtering.

6.40.3 THE VISUALISER

The display draws the ladder itself: the harmonic series as found and the same series as re-tuned, so you can see the rungs move apart or together as you turn STRETCH. It is the clearest possible picture of what the engine does - the fundamental pinned in place while everything above it slides.

6.40.4 USING IT

On a sustained synth pad, small positive STRETCH with CURVE low adds an almost imperceptible bell quality that makes the pad sound struck rather than blown. On an electric piano sound, moderate STRETCH with CURVE high exaggerates the character the instrument already has and pushes a mild Rhodes toward something far more metallic. On a bass, small negative STRETCH hollows the tone without touching the note, which can free up space in a mix that a filter sweep would not.

Two practical notes: First, Filament is at its best on material with a clear fundamental - single notes, monophonic lines, sustained pads. On dense chords, switch AUTO off and set ROOT to the key of the passage. Second, it is a genuine harmonic rebuild rather than a filter, so it responds to what is played rather than to where it sits in the frequency range: the same settings will sound consistent across the keyboard in a way that a fixed EQ curve never does.

6.41 FLANGER

A bucket-brigade flanger inspired by the Electro-Harmonix Electric Mistress, the pedal whose sound many people mean when they say flanger. Two ideas from that pedal shaped this engine. Its delay swept across an unusually wide range, roughly twenty to one, far more than most flangers and it swept logarithmically, so the comb spent equal time in every octave of its travel rather than lingering at the long end and flicking past the short. And its regeneration control, Colour, ran the feedback almost to the point of runaway, turning the comb into a resonant, metallic, jet-like tone. Both are reproduced here. The regeneration path carries the pedal's own safeguards: a high-pass so bass cannot pile up in the loop, a bucket-brigade low-pass so each pass through the loop darkens a little and a soft limiter that rounds the peaks the way the pedal's compander did, so Colour at full is loud and resonant and never runs away. The original was mono; two delay lines run here and WIDTH sets how far apart their sweeps sit. MIX and LEVEL are provided by the engine strip as they are for every engine.

RATE - Sweep speed, from a glacial 0.02 Hz to 12 Hz. The Mistress lived at the slow end; 0.1 to 0.5 Hz is its territory and the slower the sweep the more the comb reads as a moving space rather than a wobble. When SYNC is engaged the rate locks to your DAW tempo and the Division combo takes the place of this knob.

RANGE - How much of the delay span the sweep covers, which is the width of the flange. At 100 percent the comb travels from a third of a millisecond to over twelve, more than five octaves. Narrow it for a subtler shimmer that stays in one part of the spectrum.

COLOR (as in how E-H spelt it!) - Regeneration and the reason this engine exists. Below half it thickens the comb; above two-thirds it turns metallic and resonant, the notches sharpening into ringing peaks; at the top it sits on the edge of runaway. The knob is shaped so that most of its character lives in the upper third, as the pedal's did.

MANUAL - Where the sweep is centred within the delay span. With MATRIX engaged this becomes the tuning control for the fixed comb.

tone - How bright the regeneration stays. Low settings darken each pass through the loop for a warmer, more vintage resonance; high settings keep the feedback bright and glassy.

WIDTH - Stereo. Offsets the right channel's sweep from the left, from identical at 0 percent, which is effectively mono, to a quarter cycle apart at 100 percent, which puts the two combs out of step and opens the image without any cross-mixing.

MATRIX - Stops the sweep dead and leaves the comb exactly where it is, a fixed resonant filter you then tune with MANUAL and sharpen with COLOR. Engage it mid-sweep to catch a particular colour or use it with RANGE at zero as a static comb filter.

INV - Inverts the wet signal for a negative flange. The positive mix, the pedal's own, gives a peak at the low end and a full, sweeping sound; the negative mix gives a notch there instead and a hollow, more scooped tone.

SYNC + DIVISION - Locks the sweep rate to your DAW tempo. When engaged the RATE knob is replaced by the Division combo, from 32 bars down to 1/64th notes, with dotted and triplet variants.

The visualiser shows the comb itself: the frequency response of dry plus wet for each channel on a logarithmic axis, drawn from the engine's actual delay and feedback. The teeth slide as the sweep moves and sharpen as COLOR rises and a dot along the bottom rides the sweep. With MATRIX engaged the picture holds still, which makes tuning the fixed comb by eye a simple matter.

6.42 EXCITER

An exciter makes a sound brighter, livelier and clearer and makes its low end fuller, without simply turning the treble and bass controls up. It works by adding things to the sound rather than by reshaping what is already there: new harmonics above the high notes, extra weight below the bass notes and a sense of air and width at the top. Used gently, it helps a part sit forward in a mix. Used hard, it becomes a sound design tool in its own right.

The Exciter takes the best ideas from the classic hardware units that defined the category and puts them in one engine. The Exciter was inspired by the classic hardware units that defined the category, among them the Aphex Aural Exciter and Big Bottom, the BBE Sonic Maximizer and the SPL Vitalizer and by the long-standing studio technique of using a noise-reduction encoder to add air to vocals. NeuralDream's own engine brings those ideas together in one place.

Every part of the Exciter adds to an untouched copy of your sound, so when every amount is at zero you hear your sound exactly as it went in. Because the Exciter already blends its own effect with the untouched sound, you will usually want the engine's MIX control, at the right of the panel, at 100% and you set how much of the effect you hear with AMOUNT.

6.42.1 THE PANEL

The panel is arranged in four columns of controls, followed by a small live display. From left to right, the columns are LOW, MIDDLE, HIGH and MASTER, so the controls run from the bottom of the frequency range to the top, in the same order as the frequency scale of the display beside them.

6.42.2 THE LOW COLUMN

The LOW column makes the bass fuller. It takes a copy of the bass, gives it more lift when the bass is quiet and less when it is loud and adds it back, so a bass line gains weight and sustain without its peaks becoming louder.

LOW TUNE - Chooses the bass region to strengthen, from 40 Hz to 190 Hz. The default is 80 Hz. Lower settings add depth to the very bottom of the range and higher settings add body to the upper bass.

LOW AMOUNT - Sets how much bass is added, from 0% to 100%. The default is 25%. At 100% a loud bass part is lifted by around 9 dB at LOW TUNE.

LOW SHAPE - Chooses the kind of bass, from Soft at 0% to Tight at 100%. The default is 50%. Soft gives a broad, deep, rounded lift that includes the very lowest frequencies. Tight lifts the bass around and just above LOW TUNE while taking the deepest sub-bass away, which gives a dry, punchy bass that does not boom. Every position lifts the note at LOW TUNE, so turning LOW SHAPE never makes the bass disappear.

6.42.3 THE MIDDLE COLUMN

CLARITY - Adds air, from 0% to 100%. The default is 0%, which is off. CLARITY works only on frequencies above 1.2 kHz and leaves everything below 1.2 kHz completely alone. It adds a small amount of harmonic sparkle to the top of the sound and widens the stereo width that is already there: at 100% the difference between the two channels above 1.2 kHz is doubled. It does not make width out of nothing. A sound in the middle of the stereo image stays in the middle and a mono sound stays mono, gaining only the sparkle. Because it enlarges the width that is there, a part that leans towards one speaker moves a little further that way, as it would with any width control. A mono fold-down is left as it was, so nothing you do here will collapse the sound when it is played in mono.

DE-MUD - Pulls back the muddy lower midrange around 400 Hz, from 0% to 100%. The default is 0%, which is off. It removes a band centred on 400 Hz and reaching from about 200 Hz to 700 Hz, by up to 12 dB at 100%. Everything below 200 Hz and above 1 kHz stays close to where it was. Use it when a sound is thick or boxy and the bass and highs are fighting to be heard through it.

AUTO - The AUTO GAIN button, which is off by default. When it is on, the Exciter keeps its output at the same loudness as its input, measured the way broadcast loudness meters measure it (the ITU-R BS.1770 standard). An exciter can seem better simply because it makes the sound louder, so AUTO GAIN lets you judge the change in character on its own. Leave it off while you are setting the knobs, so that you hear exactly what each one does and switch it on when you want to compare fairly with the Exciter bypassed.

LISTEN - The LISTEN ADDED button, which is off by default. When it is on, you hear only what the Exciter is adding to your sound, with the untouched sound taken away. This makes it easy to hear exactly where the new harmonics, bass and air are going while you adjust the other controls. It does not change anything that is saved in a preset.

6.42.4 THE HIGH COLUMN

The HIGH column makes the top of the sound brighter and livelier. It takes the part of the sound above HIGH TUNE, adds new harmonics to it and adds the result back. The harmonics it creates are always placed above HIGH TUNE, so they add sparkle and never add mud to the lower frequencies.

HIGH TUNE - Sets the frequency above which the HIGH column works, from 550 Hz to 16 kHz. The default is 3 kHz. Lower settings brighten the upper midrange and make a part more present. Higher settings touch only the very top, for a finer, airier shine.

HIGH AMOUNT - Sets how much of the HIGH column is added, from 0% to 100%. The default is 50%. At 100% the top of the sound is lifted by around 10 dB.

CHARACTER - Chooses how the brightness is made, from 0% to 100%. The default is 50%. At the left is Clean, which creates no new harmonics at all: it brightens the quiet detail in the top of the sound much more than the loud parts, the way the Dolby A trick does, so breaths, string noise and reverb tails come forward while loud notes stay smooth. In the middle is Warm, which adds even harmonics for a rich, rounded brightness. At the right is Bright, which adds odd harmonics for a sharper, more cutting edge. Positions in between blend the neighbouring characters.

6.42.5 THE MASTER COLUMN

AMOUNT - Raises or lowers the whole effect at once, from 0% to 100%. The default is 100%. AMOUNT scales LOW AMOUNT, HIGH AMOUNT, CLARITY and DE-MUD together, so at 100% every section works exactly as its own knob is set, at 50% everything is half as strong and at 0% the Exciter adds nothing at all. It is the control to reach for when a preset is right in character but too strong or too weak for a particular mix and it is an ideal target for a macro.

DRIVE - Sets how hard the HIGH column drives its harmonic generator, from -20 dB to +6 dB. The default is 0 dB. Turning DRIVE up adds more harmonics and so more sparkle and turning it down makes the harmonics gentler. The harmonics also grow with how loud the sound is, so loud notes and attacks receive more sparkle than quiet passages, just as they do in the original hardware. DRIVE affects the Warm and Bright characters; it has no effect when CHARACTER is at Clean, which creates no harmonics.

OUTPUT - Sets the Exciter's final output level, from -12 dB to +12 dB. The default is 0 dB.

6.42.6 BUILT-IN PROTECTION

Two kinds of protection work all the time and have no controls. The first is sibilance protection. When "S" and "F" sounds dominate the top of the sound, the HIGH column automatically holds back, by up to three-quarters, so that sibilance is never exaggerated into a harsh whistle. On material without sibilant sounds, such as most synthesiser parts, it has nothing to do and stays still.

The second is protection against aliasing, which chapter 1.6 explains. The harmonic generators in the HIGH column and in CLARITY run at four times the sample rate or twice the sample rate at 88.2 kHz and above and they are designed so that they can never create anything above three times the frequency they are given. Together this means that nothing they create can fold back into the audible range as an unrelated tone, however hard they are driven.

Running at the higher rate takes a little time and so does the precise filtering that CLARITY uses to widen the sound without pushing anything to one side, so the Exciter delays the sound by 3.58 milliseconds at 44.1 kHz, 2.83 milliseconds at 48 kHz and 2.14 milliseconds at 96 kHz. NeuralDream tells your DAW about this delay and it keeps every part of the sound in step, as section 4.3 describes, so you never need to allow for it yourself.

6.42.7 THE ACTIVITY DISPLAY

The ACTIVITY display at the right of the panel shows what each column is doing to your sound right now. It runs from 20 Hz at the left to 20 kHz at the right, with the marks 100, 1k and 10k along its horizontal line. Everything above the line is being lifted and everything below it is being cut.

The red shape is the LOW column. It is the real response of the bass section at this moment, so it follows LOW TUNE and LOW AMOUNT, it changes shape as you turn LOW SHAPE from a broad, deep hump at Soft to a lift with the deepest bass cut away at Tight and it pumps as the bass section reacts to the music.

The yellow shape is the HIGH column. It rises from HIGH TUNE, its height follows HIGH AMOUNT and AMOUNT and its shape changes with CHARACTER. It moves as the Clean character reacts to the music and dips while sibilance protection holds back. The new harmonics themselves cannot be shown as a shape, because they are new sound rather than a change in level; the main visualiser shows them.

The sky-blue shapes are the MIDDLE column. A see-through sky-blue area above 1.2 kHz shows how much CLARITY is widening the top of the sound and a sky-blue dip around 400 Hz shows how deeply DE-MUD is cutting.

The two lines beneath the display are meters. The AUTO GAIN line shows how much AUTO GAIN is correcting the level, with no correction at its centre, a cut to the left and a boost to the right, up to 12 dB either way. The SIBILANCE line shows how much sibilance protection is holding back, growing from the left.

6.42.8 THE VISUALISER

The Exciter's visualiser shows the spectrum of your sound as a line and shades in everything the Exciter changes. Added bass is shown in red, added high frequencies in yellow and anything taken away in sky blue. It compares your sound before and after the Exciter but before AUTO GAIN and OUTPUT, so a simple change in level is never shown as something added. It reads both stereo channels, so the width that CLARITY adds is included. With AMOUNT at 0%, nothing is shaded.


Frequencies are spaced the way the ear hears pitch rather than in equal octaves, so the middle of the window falls at about 2 kHz and the top of the range, where the Exciter does most of its work, gets more than half the width. Every point on the curve is the average of a sixth of an octave and rises quickly but falls back slowly, like a good meter, which keeps the picture steady and smooth instead of flickering. The bass is measured over a longer stretch of sound than the highs, so low notes are drawn in fine detail while the top stays quick to react.

Three labelled lines mark LOW TUNE, the 1.2 kHz point where CLARITY begins and HIGH TUNE and they move as you turn those knobs. A faint grid marks the decibel levels and a fainter line shows the loudest level the processed sound has recently reached, falling back slowly. Across the top of the display is a WIDTH strip: the stronger the sky blue in any part of it, the wider the stereo image is at those frequencies, so you can watch CLARITY spread the top of the sound as you turn it up.

6.42.9 USING IT

To bring a dull synth pad to life, start with CHARACTER at Warm, set HIGH TUNE between 3 kHz and 5 kHz and raise HIGH AMOUNT until the pad opens up, then add a little CLARITY for air at the top and for extra width if the pad is in stereo. To make a vocal or an acoustic instrument breathe without making it harsh, turn CHARACTER to Clean, which brings up the quiet detail and leaves the loud notes alone. To give a bass more weight without more boom, raise LOW AMOUNT with LOW SHAPE towards Tight. When a busy mix sounds thick, a little DE-MUD often lets the bass and the highs through more than any amount of boost would.

Whatever you do, press LISTEN now and then to hear exactly what is being added and switch AUTO on before you judge the result against the bypassed sound, because a louder sound almost always seems like a better one at first.



7 MODULATION

7.1 HOW MODULATION WORKS

Modulation is what transforms NeuralDream from a collection of static effects into a living, breathing instrument. By connecting modulation sources to engine parameters you create movement, expression and dynamic response that evolves over time or reacts to the audio itself. NeuralDream provides twelve modulation sources organised into three types: four LFOs, four envelope followers and four macro knobs. Any of these can be routed to virtually any knob in the plugin - engine parameters, sample player controls, mix levels and more. Each routing has its own independent depth control that determines how much the modulation source affects the target parameter.

Modulation in NeuralDream works as an offset system. The APVTS (Audio Processor Value Tree State) always holds the base value of every parameter - the value you see on the knob. Modulation offsets are computed fresh every audio block and added on top of the base value then clamped to the parameter's valid range. This means modulation never corrupts saved values. When you save a preset only the base values are stored. When you load a preset or recall a DAW session the base values are restored cleanly and modulation resumes from those clean starting points. Modulation runs at block rate which means offsets are recalculated approximately every 11 milliseconds (at a typical buffer size of 512 samples at 44.1 kHz). This is fast enough for smooth sweeps and rhythmic effects though extremely rapid modulation will show stepping at this resolution.

7.2 LFOs

The four LFOs (Low Frequency Oscillators) generate repeating waveforms that cycle at a set frequency or in sync with your DAW's tempo. They are the most common modulation source for creating movement - filter sweeps, tremolo, vibrato, panning and rhythmic parameter changes. Each LFO has the following controls:

RATE - The oscillation speed. In free mode this sets the frequency in Hz. In sync mode this selects a beat division.

SHAPE - The waveform shape. Eight waveforms are available: Sine (smooth organic movement), Triangle (linear ramps up and down), Square (abrupt switching between two values), Sawtooth (linear ramp up with instant reset), Ramp (linear ramp down - the inverse of sawtooth), Sample-and-Hold (random stepped values that change at the LFO rate), Noise (continuous random output without stepping) and Random (smooth interpolation between random target values for an organic wandering movement).

SMOOTH - Output smoothing time. Applies a one-pole lowpass filter to the LFO output to soften transitions. This is particularly useful with Sample-and-Hold and Square waveforms where the abrupt value changes can cause clicks on sensitive parameters like filter cutoffs. Higher values produce a more gradual transition between values.

SYNC - Toggles between free-running mode and tempo-synced mode. When synced the LFO locks to your DAW's transport with beat divisions ranging from 1/32 notes to 4 bars.

When sync is enabled LFO phase is derived directly from the DAW's PPQ (pulses per quarter note) position. So, the LFO phase is sample-accurate, it locks precisely to the beat grid and maintains perfect alignment even after stopping and restarting transport or jumping to a different position in your arrangement. When the DAW transport is stopped the LFO continues in free-run mode so you still hear modulation effects while auditioning sounds without playback.

Each LFO can be routed to multiple parameters simultaneously. For example a single slow sine LFO could gently sweep a filter cutoff, subtly modulate a reverb decay time and slightly shift a delay feedback amount all at once creating a cohesive evolving texture from one modulation source.

7.3 ENVELOPE FOLLOWERS

The four envelope followers track the amplitude of the audio signal and convert it into a control signal. Where LFOs produce predetermined repeating patterns envelope followers respond dynamically to whatever audio is passing through the plugin - loud passages push parameters one way and quiet passages pull them back.

Each envelope follower analyses the incoming audio level and outputs a smooth control signal that rises when the audio gets louder and falls when it gets quieter. The response speed determines how closely the control signal tracks the audio dynamics - fast settings follow individual transients while slow settings capture the broader dynamic shape of phrases and passages.

Classic applications include driving a filter cutoff from an envelope follower so the filter opens on loud notes and closes between them (auto-wah), controlling a reverb mix so quiet passages get more reverb while loud passages stay dry or modulating a chorus depth so the effect intensifies during sustained notes and disappears on staccato playing.

Each envelope follower has the following controls:

Input - A source selector dropdown that chooses which audio signal the envelope follower tracks. This can be set to the plugin input or other audio sources.

SENS - Sensitivity. Controls how much the incoming audio level is amplified before the envelope detection stage. Higher values make the envelope follower respond to quieter signals.

ATK - Attack time. How quickly the control signal rises when the audio gets louder. Fast settings follow individual transients closely. Slow settings smooth over rapid level changes and capture only the broader dynamic shape.

REL - Release time. How quickly the control signal falls when the audio gets quieter. Fast settings let the envelope drop between notes. Slow settings sustain the control signal for a more gradual response.

GATE - Gate threshold. Sets a minimum audio level below which the envelope follower outputs zero. This prevents the envelope from responding to noise floor or very quiet signals, giving you a cleaner control signal with a definite off state.

7.4 MACROS

The four macro knobs are manual modulation sources - they do not generate their own movement but instead act as master controllers that you move yourself. Each macro can be routed to multiple parameters simultaneously turning a single knob turn into a coordinated multi-parameter change.

Macros are powerful for several reasons:

Performance control - Map a macro to five or six parameters and control them all with one hand during a live performance. For example a single macro could simultaneously increase delay feedback, open a filter, reduce the dry level and increase reverb mix to create a smooth transition from a clean dry sound to a washed-out ambient texture.

Preset variations - Set up a macro that morphs between two versions of the same preset. One end might be subtle and refined while the other is extreme and aggressive. The macro smoothly interpolates between the two.

DAW automation - Macros can be automated by your DAW just like any other parameter. This means you can draw automation curves that control multiple parameters simultaneously through a single automation lane.

7.5 CREATING A MODULATION ROUTING


To route a modulation source to a parameter:

Right-click any knob on a control panel. A modulation inspector overlay appears anchored near the knob, showing one row for each existing routing - source label, draggable depth track, BIP/UNI toggle and an [x] button to remove the routing - with an "+ Add source..." combo at the bottom. Click the combo and pick from the available sources (LFO 1-4, Envelope Follower 1-4, Macro 1-4) to add a new routing.

Each new routing's depth track behaves according to the source's polarity: unipolar tracks (Envelope Followers, Macros) run from 0% on the left to 100% on the right, drag right to increase depth; bipolar tracks (LFOs) place 0% at the centre with negative depth to the left and positive depth to the right. You can override the default polarity on any routing by clicking its BIP/UNI button.

Multiple sources can be routed to the same parameter. Their effects are summed together so you could have an LFO providing rhythmic movement while an envelope follower adds dynamic response on top.

To remove a routing right-click the knob again and deselect the source from the popup using the 'x' button next to the assignment.



7.5.1 UNIPOLAR AND BIPOLAR ASSIGNMENT

Each routing also has a BIP/UNI toggle that switches between bipolar and unipolar modulation. In bipolar mode (BIP) the modulation swings both above and below the knob's base value - an LFO assigned to a filter cutoff would sweep the cutoff higher and lower than the set position. In unipolar mode (UNI) the modulation only pushes in one direction from the base value. LFOs default to bipolar since their waveforms naturally swing between positive and negative. Envelope followers and macros default to unipolar since their output is inherently one-directional - an envelope follower tracks level which only goes up from zero. You can override the default for any routing by clicking the BIP/UNI button in the modulation inspector.

7.6 VISUAL FEEDBACK

NeuralDream provides extensive visual feedback for modulation activity:

A red ghost pointer on each modulated knob shows the live modulated position in real time, sliding around the knob arc as the modulation source moves. The ghost pointer is fixed-colour (a distinctive modulation red) so you can always pick out the modulated reading at a glance regardless of which engine the knob belongs to.

The animated diamonds on each routing's depth track live inside the modulation inspector overlay (right-click the knob to see them) and show each individual source's output scaled by its depth setting.

Knob badges in the top-right corner of any modulated knob show the number of active routings on that parameter; click the badge to mute or unmute all of that knob's modulation at once.

7.7 MODULATION TIPS

Start subtle - modulation depths of 10-20% often produce more musical results than extreme settings. You can always increase the depth later once you hear how the modulation interacts with the audio.

Combine sources - routing both an LFO and an envelope follower to the same parameter creates movement that is both rhythmic and dynamically responsive. The LFO provides the predictable pattern while the envelope follower adds organic variation on top.

Use sync strategically - tempo-synced LFOs create rhythmic effects that lock to your track but free-running LFOs at slightly off-beat rates create a looser, more organic feel. Try both and listen to which suits the material.

Modulate modulation - you can route a macro to an LFO's rate or depth giving you real-time control over how fast or how deep the modulation runs. This creates a second layer of expressiveness on top of the modulation itself.

Use the VIA route to control modulation with other modulators! Or use **SLOW** and **SMOOTH** to remove rough edges from Aftertouch or slow their release. You can modulate anything that has a knob inside the Synthesiser (if running NeuralDream Instrument) you have access to per note modulation as well as the GLOBAL modulation sources from the panel at the bottom of the plugin. In Effects mode, you have access to the same functionality from the GLOBAL modulators only.

8 THE VISUALISER

8.1 OVERVIEW

The visualiser occupies the top of the plugin, above the preset strip. It provides real-time visual feedback on the audio passing through NeuralDream and serves as the primary window into what each engine is doing to your sound.

The visualiser is more than a waveform display. Every one of the 40 (or so) engines has its own bespoke visualisation designed to represent the specific processing that engine performs - the Parametric EQ shows a frequency response curve, the Granular engine shows individual grain envelopes, the DARKSKY engine paints a 3D aurora and the RAVE engine streams its latent space dimensions as flowing coloured traces. Switching between engines in the layer area immediately switches the visualiser to show that engine's unique display.

8.2 VIEW MODES

Three buttons in the top corner of the visualiser select what is displayed. IN and OUT appear only in the effect version of NeuralDream; the instrument offers XFM alone.

IN - Shows the raw input waveform arriving at the plugin before any processing. This is useful for monitoring the source signal and checking levels.

OUT - Shows the final output waveform after all layers have been processed and mixed together. Comparing IN and OUT gives you an immediate visual sense of how much the processing is affecting the signal.

XFM - Shows the engine-specific visualisation for whichever transformer is currently selected in the layer area. This is the default and most informative view. Each engine's visualiser is designed to give you meaningful visual insight into what that engine is doing - not just a generic waveform but a representation of the specific processing taking place. Switch XFM off (and IN and OUT in the effect version) and the visualiser folds away to its row of buttons, so the layers below can use its space. Switch XFM back on to open it again.

The visualiser area is also where the step sequencer and Starling appear. The SEQUENCER and STARLING buttons in the preset strip switch the main visualiser to those editors and back again; they are mutually exclusive with the maximised wave display, so opening one closes the other. If the visualiser is folded away, opening the sequencer, Starling or the maximised wave display opens it and switches XFM on. Folding the visualiser away again closes whichever one is showing. See Sections 11.7 and 11.8.

8.3 FULLSCREEN MODE

Double-click anywhere on the visualiser to toggle fullscreen mode. The visualiser expands to fill the entire plugin window, hiding the layer area, engine panels and modulation strip. This gives you a large detailed view of the current visualisation which is particularly striking with engines like DARKSKY's aurora or the Neural FX mandala. In fullscreen mode you can use the left and right arrow keys to step through presets. This makes fullscreen useful for performance situations where you want a clean visual display while switching between sounds.

Double-click again or press Escape to return to the normal view.

8.4 THE DEEP SPACE BACKGROUND

Behind the visualiser sits a procedurally generated deep space scene. A domain-warped fractal Brownian motion (fBm) nebula rendered at 384x192 resolution provides a deep atmospheric backdrop with five rotating colour pallets and accents in the densest regions. 280 stars are distributed across a 3D celestial sphere with camera rotation and respond to the audio output level - louder passages cause the stars to brighten and pulse, showing off the parallax effect.

The deep space background is visible through the layer area as well since layers use a slightly transparent white background (~87% opaque) that lets the starfield subtly show through. This gives the entire plugin a sense of depth and atmosphere.

The background is hidden automatically when the visualiser enters fullscreen mode so the engine visualisation has a clean dark canvas.

8.5 ENGINE VISUALISERS

Each engine's visualiser is tailored to represent its specific processing. Here is a summary:

- ★ LP/HP Filter - Frequency response curve and slope setting, with a cutoff-frequency marker. An additional spectral overlay showing where frequency energy exists in the audio to be processed.
- ★ VCF - The filter response curve with mode indication. An additional spectral overlay showing where frequency energy exists in the audio to be processed.
- ★ Parametric EQ - A 512-point logarithmic frequency response plot with three coloured spectrum curves (red for the low band, gold for the mid band, blue for the high band) each with a gradient fill below the curve. An additional spectral overlay showing where frequency energy exists in the audio to be processed.
- ★ Noise Gate - A scrolling level history with a filled region showing gate state. A dashed threshold line with dB label, an OPEN/CLOSED text indicator and a vertical gain reduction meter on the right edge.
- ★ Delay - Animated delay taps.
- ★ Stereo Echo - A tap timeline showing left channel taps in hot pink and right channel taps in electric cyan with glow effects.
- ★ FDN Reverb - A diffusion field pattern responding to the reverb energy.
- ★ SUPERSPACE - Mode-dependent colour palette with core, mid and outer colour progressions that shift across the full colour wheel as you change algorithms.
- ★ DARKSKY - A 3D perspective aurora with flowing curtains of colour. Bright mint at the horizon rising through electric green and vivid gold to hot pink overhead. The aurora uses hardcoded vivid colours (the one deliberate exception to the Design System) for maximum visual impact.
- ★ Chorus - An animated 3D water surface with ripples driven by the chorus modulation.
- ★ Phaser - Trailing phase patterns.
- ★ Auto Pan - Panning position visualisation.
- ★ Binaural - 3D spatial position display.
- ★ Stereo Widener - A correlation meter showing the stereo relationship from mono through natural stereo to wide.
- ★ Granular - Trapezoidal grain envelope shapes showing individual grain activity in real time.
- ★ IR Convolver - A dual-layer waveform with a cyan ghost outline of the original loaded IR and a yellow live view of the processed IR after the pipeline.
- ★ Pitch Shifter - A flowing ribbon waveform with a chromatic pitch wheel showing the current effective shift amount.
- ★ Tape Simulation - The Jiles-Atherton hysteresis curve showing the S-shaped relationship between input and magnetisation.
- ★ Neural Amp - A VU-style peak meter with the loaded model name.

- ★ Neural FX (RONN) - An animated mandala pattern reflecting the neural network activity.
- ★ RAVE - Streaming latent space values displayed as flowing coloured traces, one per latent dimension.
- ★ Spectral Morph, Spectral Gate, Spectral Style Transfer, Vocoder, Cross Synthesis and Live Convolv - present a soft-brush cloud visualiser showing the carrier and modifier as spectral clouds with white needles pin-pointing the effect the affecting of the spectral transformation
- ★ Ring Modulator - A phase-space attractor plot tracing the relationship between carrier and modifier signals.
- ★ Noise Generator - A scrolling waveform coloured by the current noise type, interpolating smoothly between brown, pink, white, blue and violet as the COLOUR knob sweeps through the spectrum.
- ★ Octave Filter -- Ten vertical band columns with the active band highlighted in warm amber. Column heights represent per-band weights.
- ★ Level/Pan Engine - Ping Pong!!
- ★ BitFold - Shows a visualiser per mode demonstrating the degradation of the signal
- ★ Lustre - Banded receding reverb tail with rainbow-coloured frequency bands, per-band brush sigY ramping from low to high frequencies, vanishing point at top of panel, deep-night gradient background
- ★ Dream Delay - Three-lane cloud-particle field with periodic firings synchronised to each tap's delay period, clouds tinted blue (pitched up) or pink (pitched down) based on each tap's shift amount captured at spawn
- ★ Aperture - Vertical slice of frequencies included and excluded by the 'brickwall'
- ★ Particles - a three-dimensional box in which every living particle is a point of light with a motion trail. Left and right is its position in the stereo field, up is its pitch and depth is its age: a particle is born at the back, at the emitter and flies towards you as it sounds, dying at the front. First-generation particles are cyan, children amber and grandchildren magenta and the camera sways slowly so the depth reads even when nothing is changing
- ★ Exciter - The spectrum of your sound as a line, with everything the Exciter changes shaded in: added bass in red, added highs in yellow and anything taken away in sky blue. Frequencies are spaced the way the ear hears pitch, so the top of the range gets more than half the width and the curves are smoothed so the picture stays steady. Labelled markers follow LOW TUNE, the 1.2 kHz point where CLARITY begins and HIGH TUNE and a strip across the top shows how wide the stereo image is at each frequency. It compares the sound before AUTO GAIN and OUTPUT, so a change in level is never shown as something added. The panel also has its own small ACTIVITY display, which section 6.42.7 describes.

9 UNDO AND REDO

9.1 OVERVIEW

NeuralDream maintains a 12-level undo history that lets you step backward and forward through recent changes. This is a safety net for experimentation - you can make bold changes to your processing chain knowing that you can always step back to where you were.

9.2 WHAT IS CAPTURED

The undo system captures six types of action:

Parameter changes - Moving any knob or changing any value in the plugin.

Engine additions - Adding a new transformer engine to a layer's chain.

Engine removals - Deleting an engine from the chain.

Engine reordering - Dragging engines to a new position in the chain.

Engine swaps - Using the triple dot menu to exchange one engine for another

Preset loads - Loading a preset replaces the entire plugin state. Undo restores the state you had before loading.

Modulation routing changes - Adding, removing or modifying modulation assignments and depths.

The history is stored as a circular buffer of complete state snapshots using the memento pattern. When you undo the plugin restores the previous snapshot in its entirety - every parameter, every engine, every routing returns to exactly where it was. Redo steps forward again through snapshots you have undone.

9.3 USING UNDO AND REDO

The undo and redo buttons are located in the preset strip to the right of the SAVE button. Click the left arrow to undo and the right arrow to redo. The buttons are greyed out when there is nothing to undo or redo.

If you undo several steps and then make a new change the redo history is discarded from that point forward - the new change becomes the latest entry in the history. This is the same behaviour you would expect from undo in any application.

The undo history is not saved with presets or DAW sessions. It exists only for the current plugin session and is cleared when the plugin is removed from the DAW or when you press INIT or CLR.

10 CREATIVE IDEAS/TIPS

10.1 GETTING STARTED

If you are new to NeuralDream the simplest way to begin is to load a factory preset and explore what it does. Select a preset from the browser then click through the engines in the layer area to see what each one is doing. Adjust knobs, toggle engines on and off and listen to how each stage contributes to the overall sound. The undo system means you can experiment freely without risk.

When you are ready to build from scratch press INIT for a clean slate. Add a single engine to Layer 1 - a Chorus or Delay is a good starting point - and shape the sound from there. Add a second engine after it and listen to how the order changes things. A Chorus into a Delay sounds very different from a Delay into a Chorus. This kind of exploration is where NeuralDream rewards curiosity.

10.2 LAYER STACKING

One of the most powerful techniques in NeuralDream is running the same input through multiple layers with different processing on each then blending the results with the layer mix and gain controls.

Try setting up two layers both drawing from Plugin Input. On Layer 1 add a Pitch Shifter set to +7 semitones (a perfect fifth) with a low mix. On Layer 2 add a Stereo Echo with long feedback. The result is a harmonised delay - the echoes carry the pitch-shifted version while the dry sound stays clean. Adjust each layer's gain independently to find the right balance.

For ambient work try three layers: one with a short FDN Reverb for early reflections, one with a DARKSKY for the long tail and one with a Granular engine set to high density and slow speed for a frozen texture layer underneath. Each layer contributes a different dimension to the space.

10.3 SERIAL LAYER ROUTING

Setting a layer's source to the output of another layer creates a serial processing chain. This is useful when you need more engines than a single chain can comfortably manage or when you want to control the blend at an intermediate point.

For example set up Layer 1 with a Tape Simulation and Parametric EQ processing the plugin input. Then set Layer 2's source to Layer 1 and add a Granular engine followed by a DARKSKY reverb. Layer 2 is now granulating and reverberating the tape-processed signal. Adjust Layer 1's gain to control how much level reaches Layer 2 without affecting Layer 1's own contribution to the final mix.

Remember that inter-layer routing taps the audio before the output level control so you can fade Layer 1 out of the final mix entirely while still feeding it into Layer 2 at full level. This lets you use layers purely as processing stages that are heard only through downstream layers.

10.4 THE NEURAL ENGINES

The three neural engines each offer a fundamentally different creative approach.

Neural FX (RONN) is a discovery tool. Every seed number produces a different neural network with different sonic characteristics. Some seeds produce subtle warming, others create aggressive distortions and others generate entirely alien textures. There is no way to predict what a seed will sound like - the only way to find out is to try it. Spend time sweeping through seeds with audio running and save the ones you like as sub-presets. Building up a library of favourite seeds is part of the creative workflow.

Neural Amp is the most predictable of the three. Load a NAM model of a specific amplifier or pedal and you get a faithful reproduction of that hardware's character. The creative possibilities open up when you place it in unexpected contexts - try a guitar amp model on a vocal, a bass amp on a synth pad or a tape machine model on a drum bus. The INPUT knob lets you drive the model harder or softer than the original hardware was designed for which often produces interesting results at the extremes.

RAVE is the most transformative. Because it encodes audio into a learned latent space and decodes it back through the lens of its training data the output always carries the timbral DNA of whatever the model was trained on. A model trained on violin recordings will impose string-like qualities on any input. A model trained on speech will make everything sound vaguely vocal. The creative power comes from the collision between your input audio and the model's learned characteristics - the further apart they are the more surprising the results.

10.5 SPECTRAL PROCESSING

The dual source spectral engines (Spectral Morph, Spectral Gate, Vocoder, Live Convolve, Spectral Style Transfer) all work by analysing two audio signals in the frequency domain and combining them in different ways. The key creative decision with all of them is choosing what to use as the modifier signal.

The modifier can come from a dedicated sample player, the plugin input or the output of another layer. Loading a rhythmic loop as the modifier and feeding a sustained pad as the carrier through the Vocoder imposes the rhythm of the loop onto the tonal character of the pad. Swapping the two - rhythmic carrier with sustained modifier - produces a completely different result from the same two signals.

Spectral Style Transfer is particularly interesting for sound design. Load a recording of rain as the modifier and play a synth through the carrier - the synth takes on the spectral character of rainfall while retaining its own pitch and rhythm. The FREEZE button captures the modifier's spectral envelope at a single moment and holds it, letting you continue processing the carrier against a frozen snapshot. This is useful for capturing a specific timbral quality and applying it to different material.

10.6 MODULATION AS A CREATIVE TOOL

Modulation can turn a static preset into something that evolves and breathes over time. Here are some approaches to try.

Slow LFO on reverb decay - A sine LFO at 0.1 Hz with modest depth on a reverb's decay time makes the space gently expand and contract creating a breathing quality that works beautifully under ambient material.

Envelope follower on filter cutoff - Route an envelope follower to the VCF cutoff for a classic auto-wah effect. The filter opens when you play loud and closes when you play soft. Try this with the VCF set to band pass mode for a more vocal quality.

Macro-controlled wet/dry - Route a macro to the MIX controls of several engines simultaneously. At one extreme the sound is completely dry. At the other extreme every effect is at full wet. The macro gives you a single knob that crossfades between clean and processed.

Sample-and-hold on pitch - Set an LFO to sample-and-hold waveform, sync it to 1/16 notes and route it to the Pitch Shifter's SHIFT parameter with the SCALE set to a musical key. The result is randomised melodic pitch sequences that stay in key - a generative melody machine.

Cross-modulation - Route an LFO to the rate of the Phaser's internal sweep. The phaser speed now rises and falls creating an organic, non-repeating sweep pattern that is more interesting than a fixed rate.

10.7 WORKING WITH SAMPLES

The sample player is far more than a simple loop trigger. Here are some creative ways to use it.

Position mode as a sound source - Set the playback mode to Position and assign an LFO to the POSITION knob. The LFO now scans through the sample at whatever speed and pattern the LFO waveform dictates. A slow triangle wave gently scrubs back and forth. A fast sample-and-hold jumps to random positions creating glitchy, granular textures. A tempo-synced sawtooth produces a rhythmic retriggering effect.

Layered samples as modifiers - Load different samples into the modifier slots of dual source engines across multiple layers. A vocal sample driving a Vocoder on Layer 1, a drum loop driving a Spectral Gate on Layer 2 and a texture driving a Ring Modulator on Layer 3 - each layer responds to a different rhythmic and tonal character.

Time stretch for tempo matching - Enable stretch on a sample and NeuralDream automatically adjusts its playback speed to match your DAW's tempo. This lets you use rhythmic loops and phrases at any tempo without pitch artefacts. Combine this with the SHIFT knob for independent pitch transposition - slow a loop down to half speed while shifting it up an octave for example.

10.8 BUILDING PRESETS FOR PERFORMANCE

If you are using NeuralDream in a live context consider building your presets with performance in mind.

Set up macros as your primary performance controls. Map Macro 1 to effect intensity (mix levels, feedback amounts, modulation depths), Macro 2 to tonal character (filter cutoffs, EQ gains, drive amounts), Macro 3 to spatial width (reverb mix, stereo widening, panning depth) and Macro 4 to chaos (granular density, neural FX depth, pitch shift range). Four knobs give you comprehensive real-time control over the entire processing chain.

Use the preset browser's category system to organise presets by song or set rather than by sound type. Name them clearly so you can find them instantly under stage lighting.

Remember that NeuralDream's preset system saves the complete state including sample audio embedded in the preset file. This means your live presets are fully self-contained - you do not need to worry about missing sample files on a different machine.

10.9 MAKE YOUR OWN PRESETS

The factory presets are starting points not destinations. Every parameter in NeuralDream can be automated, modulated and saved. Experiment freely - press INIT to return to a clean slate whenever you want to start over. When you create something you love hit SAVE, choose a category, write a note describing what inspired it and it becomes part of your personal library.

The best presets often come from accidents. Route a modulation source to the wrong parameter, load an unexpected sample as a modifier or push a feedback control further than you intended. If it sounds interesting save it immediately - you can always refine it later but the moment of discovery is hard to recreate.

11. SYNTHESISER

11.1 THE SYNTHESISER ENGINE

The synthesiser is a polyphonic wavetable sound generator that lives inside a layer as an audio source, in exactly the same way as the sample player. Where the sample player reads a recording from disk, the synthesiser generates its own audio from oscillators driven by the MIDI notes you play. Everything downstream is unchanged: the synthesiser feeds the layer's transformer chain, its output level and pan are set by the usual layer source controls and the whole layer behaves like any other.

To use it, set a layer's source to Synth from the source selector. In NeuralDream Instrument a new instance and the INIT state both start with the synthesiser already loaded on Layer 1, so a fresh track is immediately playable. You may run the synthesiser into any chain of engines you like, modulate it from the global modulation strip and save the result as a preset. Because it is a layer source rather than a fixed top-and-tail instrument, you can also layer two synthesiser voices by giving two layers a synth source or feed one synthesiser through another layer's effects.

AT A GLANCE:

Up to thirty-two voices of polyphony • three multi-mode oscillators • wavetable, granular, partials or modal - each with per-oscillator warping (PD, Fold and GRIT) • drag-and-drop sample loading into the granular and partials engines • a physically modelled exciter and bucket-brigade resonator in modal mode, with optional sidechain excitation • a noise source • a four-channel mixer with constant-power panning • two multimode filters • dedicated amplitude and filter envelopes - two per-voice LFOs • a six-source per-voice modulation matrix with VIA depth control • full MPE and Poly Aftertouch with per-note pitch bend, pressure and slide • a global voice cap that persists across instances.

11.3 SYNTHESISER ARCHITECTURE

The signal flows from three oscillators and a noise source, through a four-channel mixer, into two filters and finally the amplifier - with two envelopes and two LFOs shaping the sound and a voice allocated to each note you play. The sections below follow that path.

11.3.1 VOICES AND POLYPHONY

The synthesiser provides up to thirty-two voices. Each note you play is assigned its own voice and each voice is a complete copy of the whole signal path - its own oscillators, filters, envelopes and per-voice LFOs - so notes evolve independently of one another rather than sharing one set of movements. How many voices may sound at once is set by the Voice Limit, described in Section 11.6: the default is sixteen and a preset can choose anything from a single voice up to thirty-two. When every allowed voice is in use, the oldest is reused for the next note. The sustain pedal is supported, holding voices through note-off until the pedal is released. Note that UNISON multiplies what each note costs: a note played with UNISON at 4 uses four voices from the same pool.

Analog introduces a small, controlled amount of voice-to-voice variance, subtly detuning and varying each voice so that stacked notes and unison-like textures have the gentle movement of analogue hardware rather than the perfect uniformity of digital oscillators. At zero the synthesiser is mathematically precise; raise it for life and width.

11.3.2 THE OSCILLATORS

There are three identical oscillators, OSC 1, OSC 2 and OSC 3. Each of the three oscillators can run in one of four modes, chosen by the small button in the top-right of each OSC panel. Press the button to cycle through the modes in turn. Select WT for Wavetable Synthesis, GR for the Granular Engine, PA for Partial Resynthesis or MO for the Modal Engine:

WT: The classic mode described above: the oscillator reads a single-cycle table at the pitch of the note, with Position sweeping the timbre. This is the default and the only mode that existed before this version.

GR: The oscillator plays back a loaded audio sample as a cloud of short overlapping grains, giving evolving, textural and time-stretched results from any recording.

PA: The oscillator resynthesises a loaded audio sample from its harmonic content - a bank of tuned sine partials whose levels follow the sample - for smooth, additive, keyboard-tracked tones.

MO: The oscillator becomes a struck object. A built-in exciter - an oscillator, a noise source or your DAW sidechain - is fired into a resonator that rings at the pitch of the note, producing plucked, bowed, blown and struck sounds that are generated rather than sampled. Modal mode needs no sample and no wavetable, so nothing is loaded onto it and the wave display is not used.

Any audio file can be dropped directly onto the wave-display area of an oscillator running in Wavetable, Granular or Partial mode. Modal mode does not read audio, so it does not accept dropped files. A dialog appears - Convert dropped sample to: **GRAINS / PARTIALS / WAVETABLE** - with an X in the corner to cancel. Beneath the three choices sit two switches, **TRIM START** and **TRIM TAIL**, both off by default. They remove leading and trailing silence from the recording before it is handed to any of the three converters and they are described fully in Section 11.3.2.5. The first two options load the file as the source for Granular or Partial mode, switching the oscillator accordingly; dropping a new file onto an oscillator already in GR or PA mode replaces its source. The third option converts the recording into a scannable wavetable. The file is copied into your user wavetable library (under an Imported category, transcoded to WAV if necessary) and handed to NeuralDream's resynthesis engine, which first classifies it: a genuine single-cycle wavetable - one whose length is an exact multiple of the 2048-sample frame, whose frames measure close to one cycle each or which carries a Serum 'clm' chunk declaring its frame size - is loaded directly, frames intact. Anything else is treated as an ordinary recording and rebuilt by harmonic resynthesis into a position-indexed wavetable: the audio is analysed in overlapping windows and each window resynthesised as one 2048-sample single cycle, producing up to 1024 frames with an adaptive analysis hop (never finer than 6 ms) so the frame count scales to the length of the file and POS sweeps the whole recording start to end. Left and right channels are analysed and baked independently, with a single shared phase-alignment rotation per frame that keeps the stereo image coherent and stops adjacent frames fighting as POS scans. The root pitch is measured from the audio itself by harmonic product spectrum, with the filename used only as a witness. The result behaves exactly like any library wavetable: a static POS is a frozen single cycle at correct pitch, a swept POS scans the timbre of the recording and the 3D view shows its true terrain.

Root note and pitch - When a sample is loaded, NeuralDream sets its root from the filename: a note name appearing as its own word - for example the 'F' in 'Auri_F_Grains_Ambient.WAV' or the 'C2' in 'Cello_C2.wav' - becomes the root, with the octave taken from the name where given and defaulting to the middle octave where not. If the filename carries no note, the root is middle C (C3). The rule is deliberate: pitch detection is meaningless on unpitched material such as drum loops, so the root is always predictable - C3 plays the sample exactly as recorded and other notes transpose around it. Tag pitched samples with their note in the filename to have them land in tune on the keyboard automatically.

Stereo throughout - Samples are analysed and played back in full stereo - in Granular mode the grains preserve the source's stereo image and in Partial mode the left and right channels are analysed into independent partial sets (a mono file simply shares one set across both). Quiet third-party samples are peak-normalised on load, using a single shared gain across both channels so the stereo balance is preserved while a quiet recording is brought up to a usable level.

Warping (PD, Fold, GRIT) and the mixer, filter, envelope and modulation sections all work identically whatever mode an oscillator is in - the mode only changes how the raw tone is generated.

The controls common to all three are:

Octave Coarse - tuning in octaves.

Frequency Coarse - tuning in semitones, relative to the played note.

Fine - Fine tuning in cents, for detuning oscillators against each other.

Position - In Wavetable mode, the read position through the wavetable - the timbre control. On the eight simple shapes (Sine, Saw, Ramp, Triangle, Square, Pulse, Trapezoid and Parabolic) it sets the width of the shape instead, as described under Width of the simple shapes in Section 11.3.2.1. In Granular mode it is the read position in the loaded sample and in Partial mode it scrubs through the sample's analysis frames; in both, sweeping POS scans through the sample while the pitch stays fixed to the played note.

Level / Pan - The oscillator's contribution to the mixer and its placement in the stereo field.

Each oscillator also has a set of specialise buttons in the Wave Display on the right-hand side of the the plugin, e.g. Wavetables have buttons for further processing of the Wavetable and Audio, namely WARP and GRIT, while rabUlat has GRAINS and Partial has PARTIALS and PROCESS. many of these additions are covered in full in Section 5.

If you are not using the Wavetable Mode in OSC1 and/or OSC2, you will sometimes see that the FM and SYNC options described below become greyed out and unavailable:

FM Amount - sets how strongly OSC 2 modulates OSC 1. OSC 1 is the carrier and OSC 2 is the modulator and the modulation is applied to OSC 1's read phase, so as you raise the amount OSC 1 gains sidebands and grows brighter and more metallic. With wavetables as the operators this ranges from subtle motion to clangorous, bell-like timbres.

Sync 2 to 1 - hard-syncs OSC 2 to OSC 1: every time OSC 1 completes a cycle it forces OSC 2 to restart its own. Detuning OSC 2 upward then produces the classic hard-sync sweep, a vocal, tearing timbre that responds strongly to OSC 2's pitch.

11.3.2.1 WAVETABLE SYNTHESIS

A short grounding in the technique helps the controls make sense. If you already know wavetable synthesis well.

Every sound is, at heart, a repeating shape. A single repetition of that shape is one cycle of the waveform and a classic oscillator simply reads one fixed shape - a sine, a sawtooth, a square - over and over at the pitch of the note. The timbre is whatever harmonic content that one shape contains and it never changes.

A wavetable takes that idea further. Instead of one shape, it stores a whole sequence of single-cycle shapes, one after another, like the frames of a short film. Each frame is a complete cycle in its own right and the frames are arranged so that moving from one to the next changes the tone in a musically useful way - perhaps from a soft sine through increasingly hollow and reedy shapes to a bright buzz. The oscillator reads one position in that sequence at a time, so by choosing or sweeping the read position you choose the timbre.

The Position control is what selects that read point. Rather than jumping abruptly from frame to frame, the oscillator interpolates smoothly between neighbouring frames, so moving Position glides the tone continuously through the table.

Automate or modulate Position and the timbre moves on its own - this is the signature wavetable sound, from slow evolving pads to fast vowel-like motion.

Reading a stored shape at an arbitrary pitch raises a practical problem: aliasing. If you play a bright, harmonically rich frame very high, harmonics that sit above half the sample rate fold back down as inharmonic, metallic tones. To prevent this, each table is pre-processed into a set of band-limited versions - progressively gentler copies with the too-high harmonics removed - and the oscillator picks the right one for the pitch being played. This work is done away from the audio path, so it never interrupts playback and the result is a clean tone across the whole keyboard.

The wavetable library is where the available tables live. They are organised as a browsable folder tree - typically pack, then category, then individual table - and are loaded from two places: a master collection installed with the plugin and a user folder you can add your own tables to. The browser shows this hierarchy directly and a SIMPLE folder at the top offers eight basic single-cycle waves (sine, sawtooth, ramp, triangle, square, pulse, trapezoid and parabolic) for when a single classic waveform is all you need. A selected table is remembered by its location in the library, not by a position number, so adding or removing other tables never disturbs a sound you have already built.

If you add new wavetable files to the user folder while the plugin is open, press SCAN in the preset browser. As well as refreshing the preset lists, SCAN now rebuilds the wavetable library from disk so the new files appear in the browser and it does so without disturbing the tables your oscillators are currently using.

Wavetable is the oscillator's default mode and the one the other two build on. In this mode the oscillator generates its tone from a wavetable - a stored sequence of single-cycle waveforms, arranged like the frames of a short film, where each frame is a complete waveshape and moving from one to the next changes the timbre. The oscillator reads one position in that sequence at the pitch of the note, so the pitch comes from the keyboard and the *tone* comes from where you are reading in the table.

Because the shapes are stored and read at any pitch, the mode is anti-aliased automatically: each table is pre-processed into band-limited versions with the too-high harmonics removed and the oscillator picks the right one for the note being played, so the tone stays clean from the lowest key to the highest without any action on your part.

Choosing a wave. The Wave control selects what the oscillator reads. This can be one of eight basic single-cycle shapes - sine, sawtooth, ramp, triangle, square, pulse, trapezoid and parabolic, grouped in the SIMPLE folder - for when a single classic waveform is all you need or any table from the wavetable library, browsed as a folder tree of packs, categories and individual tables. A selected table is remembered by its place in the library rather than by a number, so adding or

removing other tables never disturbs a sound you have already built. New tables added to the user folder appear after pressing SCAN in the preset browser.

Width of the simple shapes. On the eight simple shapes the POS knob sets the width of the shape. This is what pulse-width modulation does to a square wave, applied to every simple shape. At POS 0% the shape is complete. Turning POS up squeezes one whole cycle of the shape into a shorter part of each period and holds the shape's final level for the rest of the period.

On the square this is classic pulse width: POS 50% gives a 25% pulse and POS 99% gives a very thin one. The pulse shape narrows in the same way from its own 25%. On the saw and the ramp the cycle becomes a steeper and steeper spike, on the sine it becomes a short burst and the triangle, trapezoid and parabolic become narrower versions of themselves. At POS 100% the width is zero and the oscillator is silent.

Every squeezed shape is kept centred on zero, as an analogue oscillator's output is. A shape therefore gets thinner and quieter as it narrows instead of becoming lopsided. It is anti-aliased like everything else in the oscillator and stays clean at every width and every pitch. Routing an LFO to POS gives classic pulse-width modulation on the square and the same kind of movement on the other shapes. The wave display draws the shape at its current width. **The healed-table star:** Some tables arrive with a small mismatch between their last frame and their first – the two ends of the Position range do not quite meet. Sweep such a table cyclically, so the read position passes from the end back to the beginning and those two frames are spliced together mid-cycle: the mismatch sounds as a click on every pass. On load, NeuralDream measures the seam between the final frame and the first and, whenever it is large enough to click, gently crossfades the last few frames into the first so the table becomes seamless. The body of the table – its character, fuzzy, buzzy or otherwise – is never altered, the first frame is never altered and the adjustment is confined to a handful of closing frames whose step-to-step change stays within the table's own natural motion. This applies to every table alike: factory, imported or added by hand. A table that has been healed shows an asterisk after its name in the wave display's selector – a long name may be shortened to fit but the star always survives at the end. The star is purely informative: nothing needs doing and tables whose ends already meet or whose seam is too small to matter, load untouched and unstarred.

Sweeping the timbre. The Position (POS) control chooses the read point in the table. Rather than jumping from frame to frame it interpolates smoothly between neighbouring frames, so moving Position glides the tone continuously through the table – from, say, a soft sine through increasingly hollow and reedy shapes to a bright buzz. The eight basic single-cycle waves have only one frame. On those, Position sets the width of the shape instead, as described under Width of the simple shapes above. Automate or modulate Position and the timbre moves on its own – the signature wavetable sound, from slow evolving pads to fast, vowel-like motion.

Shaping the wave. Everything a wavetable oscillator plays can be reshaped by the warping engines – PD and Fold for harmonic reshaping and GRIT for saturation and lo-fi – covered under Wavetable Warping below. Warping applies to the basic waves as readily as to library tables and because each oscillator warps independently you can keep one oscillator clean as a foundation while another is heavily shaped.

The tuning, level, pan, mixer, filter, envelope and modulation controls are common to all three oscillator modes and are described in the sections around this one; Wavetable mode simply determines that the raw tone is read from a stored waveform. To play a recorded sample instead of a stored waveform, switch the oscillator to Granular or Partial mode with the WT | GR | PA toggle (Sections 11.3.2.2 and 11.3.2.3).

WAVETABLE WARPING

Warping reshapes an oscillator's tone before it reaches the mixer. It applies to whatever the oscillator is playing - a basic wave, a library wavetable, a grain cloud or a partial spectrum - and every warp control is a modulation target, so the shaping can move under an envelope, LFO or expression. There are two warping engines per oscillator: PD and Fold (harmonic reshaping), reached from the WARP overlay and GRIT (saturation and lo-fi), reached from the GRIT overlay. Each oscillator has its own independent set of both.

This section is deliberately the last in the synthesiser documentation. Warping is an area we intend to extend and keeping it at the end means new warping features can be added here as further sub-sections without renumbering the rest of the manual.

PD - PHASE DISTORTION

PD applies phase distortion in the style of the classic Casio CZ synthesisers. Rather than changing the stored shape, it changes the rate at which the oscillator reads through one cycle: the read position is sped up through part of the cycle and slowed through the rest. The waveform is read faster in one region and dwells longer in another, which concentrates energy and adds harmonics - a sine gradually takes on a sharper, more resonant, almost filter-swept character as the amount is raised. The range is 0 to 100%, where 0 is the untouched waveform. Because it acts on read phase rather than amplitude, PD keeps a smooth, vocal quality that pairs well with wavetables.

FOLD - WAVEFOLDING

Fold is a wave-folder. As its drive is raised, peaks that would simply clip in a normal saturator are instead reflected - folded - back on themselves, each fold adding a new set of harmonics. The result grows brighter and more complex in a way that stays musical rather than harsh and it is especially effective on smooth source waves such as sines and triangles, which have little harmonic content of their own for the folder to work with. The range is 0 to 100%, where 0 leaves the waveform unfolded.

FOLD SYM - FOLD SYMMETRY

Fold Sym sets the symmetry of the fold by applying a DC bias to the signal before it reaches the folder. With no bias the folding is symmetrical about the centre and produces mainly odd harmonics; biasing the signal off-centre makes the folding asymmetrical and brings in even harmonics, shifting the tone and adding a vowel-like or hollow colour. The control is bipolar, from -100% through 0 (symmetrical) to +100%, so you can push the bias either way. It is at its most expressive as a modulation target - move it with an LFO or with Slide for a living, shifting timbre.

FOLD TYPE

A dropdown selecting the transfer function of the folder itself - the mathematical curve the driven signal is folded through. Sine and Triangle are the two classic shapes: they differ in how sharply they reflect the signal at each fold, one rounder and gentler, the other harder and brighter. Beneath them, the menu lists the other oscillators - OSC 2 and OSC 3 when OSC 1 is being shaped and so on - whenever those oscillators are currently in Wavetable mode and choosing one does something rather special: that oscillator's live wavetable frame becomes the folding curve.

The technique is shapetable distortion, a direct descendant of digital waveshaping synthesis (Le Brun, "Digital Waveshaping Synthesis," Journal of the Audio Engineering Society, 1979), with the folder's drive pushing the signal through - and, at high drive, repeatedly around - one full cycle of the source table, each pass through the curve adding the harmonic signature of that table's shape.

What makes the oscillator sources musically distinct from the fixed shapes is that they move. The transfer curve tracks the source oscillator's Position control, so as the source morphs through its table the character of the distortion itself evolves - put a slow LFO on the source's Position and the fold breathes; sweep it and the shaped oscillator's spectrum is re-drawn in real time. The source oscillator does not need to be audible: turn its mixer channel to zero and it serves purely as a shaping curve, a dedicated "distortion oscillator." The wave display shows the truth throughout - the shaped oscillator's 2D cycle and 3D terrain both render the folded result and follow the source's motion live. One design note: the transfer uses the source's global Position; per-voice Position modulation of the source is not reflected in the curve, a deliberate trade that keeps the shaping mathematics off the per-voice hot path.

Anti-aliasing is preserved for arbitrary tables by the antiderivative method (first-order ADAA - Parker, Zavalishin and Le Bivic, "Reducing the Aliasing of Nonlinear Waveshaping Using Continuous-Time Convolution," Proc. DAFx-16): the engine maintains a numerically exact periodic antiderivative of the source frame, rebuilt only when the source's table or Position actually changes, so folding through a table costs no more aliasing than folding through the built-in sine. Two practical rules complete the picture. Only oscillators currently in Wavetable mode are offered - a Granular or Partials oscillator has no single frame to fold through, so it simply does not appear in the list. And if a preset recalls a source that has since left Wavetable mode, the selection is remembered and shown but the folder falls back gracefully to the Sine shape until the source returns. Saved presets from earlier versions are untouched: the original Sine/Triangle setting is stored exactly as before and the oscillator-source choice lives in a separate, additive parameter.

GRIT - Saturation and Lo-Fi

GRIT is the second warping engine, sitting alongside PD and Fold and reached from the GRIT overlay on each oscillator. Where PD and Fold reshape the waveform harmonically, GRIT degrades and drives it - a per-oscillator saturation and lo-fi stage that adds weight, dirt and vintage-digital character before the signal reaches the mixer. Like every warp control, each of its four controls is an independent modulation target.

DRIVE (0-100%) - saturation drive. As it is raised the waveform is progressively overdriven, adding harmonics and thickening the tone from a gentle warmth to a hard, distorted edge.

CHAR (0-100%) - character. Introduces an even-harmonic asymmetry to the saturation, shifting the distortion from a symmetrical, odd-harmonic bite toward a richer, more vintage even-harmonic colour.

BITS (1-16 bit, 16 = clean) - bit-depth reduction. At 16 bits the signal is untouched; lower it and quantisation noise and grit build up, from subtle crunch down to heavily crushed, low-resolution tones.

RATE (0-100%, 0 = full rate) - sample-rate reduction. At 0 the signal runs at full rate; raise it to decimate the sample rate, adding aliasing and the harsh, ringing, digital character of downsampling.

The four combine freely and reward modulation: a little DRIVE with a touch of CHAR thickens a clean oscillator; BITS and RATE together give a lo-fi, degraded voice; and because GRIT is per-oscillator you can keep one oscillator clean while another is heavily gritted. Try routing an envelope to DRIVE so notes bite harder on their attack or Pressure to BITS so leaning into a held note crushes it further.

11.3.2.2 GRANULAR MODE (GR)

In Granular mode the oscillator plays the loaded sample as a cloud of overlapping grains. Where the grains read from, how long they are, how many overlap and how far they scatter are all under your control and, being ordinary oscillator controls, every one is a modulation target, so the whole cloud can move under an envelope, LFO or expression, per note. The four core controls sit in a 2×2 grid, with the GRAIN SYNC button beside SIZE.

The engine underneath deserves a word, because its design is what makes the cloud smooth. Grain length and grain timing are governed by two independent clocks. SIZE sets how long each grain lives; a separate onset clock decides when new grains start, at an interval of SIZE divided by the number of GRAINS - so GRAINS is really an overlap factor, in line with the standard treatment of granular scheduling in the DSP literature (overlapping smooth envelopes must sum coherently or the overlap itself prints as amplitude modulation - the classic constant-overlap-add condition). Divorcing the two clocks removes the periodicity that a death-triggers-birth scheduler bakes in at 1/SIZE. On pitched material the onset interval is additionally snapped to whole periods of the played note, so the overlapped copies add in phase rather than combing; and each free-mode grain is scaled to the source's reference level (a 256-point equal-energy probe), so a decaying one-shot doesn't print its own envelope at the grain rate. The result, measured with a 3–30 Hz amplitude-modulation metric against the engine's own renders, is a held note that sits still.

GRAINS - the number of grains sounding at once, from 1 to 100. In free (unsynched) mode this is the overlap of the stream: at 1 a single full-SIZE grain plays and is seamlessly replaced when it ends - bar-long grains are perfectly usable - while higher values stack more, shorter-spaced grains for thicker textures. With GRAIN SYNC engaged it is the number of grains fired in the burst at each division.

SIZE - grain length, set continuously from about 3 ms to 30 seconds. Short grains emphasise the grain rate itself; long grains preserve the source's character and movement. The top of the range turns the engine into something else again: at large sizes with GRAINS at 1, a single grain plays the whole recording - or as much of it as exists - and is seamlessly replaced when it ends, so the granular oscillator becomes, in effect, another way to play a sample, with the full polyphonic voice architecture, envelopes, modulation and warping attached and no need to touch SPREAD or JITTER at all. The readout shows milliseconds below a second and seconds above; when spawning is tempo-synced it adds the musical length in brackets - for example 1.00 s (1/2 bar) - from the host tempo and time signature.

SPREAD - the width of the region grains are read from, as a percentage of the whole sample, centred on POS. At 0% every grain reads exactly the POS point - the classic granular freeze; raising it widens the region symmetrically for a more diffuse cloud. The knob's response is tapered so that half its sweep covers 0–10%, with fine fractional values (0.1%, 0.5%...) available at the bottom where the musically critical territory lies.

JITTER - timing randomness, with a role that follows the spawning mode. In free mode it randomises the onset clock's period: each interval is drawn as the nominal value scaled by ±JITTER, so 0% is a metronomic weave, 100% scatters intervals anywhere up to double the nominal and the average density is preserved. In GRAIN SYNC mode it delays each grain of the burst after the division boundary by a random amount up to JITTER percent of the grain's own length. Where SPREAD scatters where grains read from, JITTER scatters when they fire. At 0% there is no hidden timing randomness of any kind.

POS - the read position in the loaded sample, start to end. Held still (with SPREAD at 0%) it freezes one moment; swept, it scans the read point through the recording without changing pitch - the granular equivalent of scrubbing a tape by hand. POS is an anchor: grains never walk the sample by themselves, in any mode, so what you set (or modulate) is exactly what plays. A clock-synced sweep (an LFO in Sync mode running a saw over a bar) turns a loop into a tempo-locked, re-triggerable scan. Position wins at the sample's end: a grain asked to start near the end starts exactly there and simply plays what room remains, rather than sliding earlier to fit.

GRAIN SYNC - tempo-locked spawning. With it off, the free two-clock engine described above runs from note-on. Click it to choose a musical division and a burst of GRAINS grains fires at every division boundary, sample-accurately, for strict rhythmic slicing. Synced grains are deliberately not level-normalised: each slice carries the true dynamics of the audio under it, so a quarter-note slicer keeps the light and shade of the source - flattening it would destroy exactly what the technique trades on. When a quantised modulation drives POS in this mode (a sample-and-hold stepping through slice positions), the quantiser uses beat-start semantics: a 100% unipolar source with four steps lands on 0, 25, 50 and 75% - the starts of the four quarters - never on 100%.

Grain edges and the tape-splice rule - Each grain is shaped by a square-by-default envelope with a minimum 10 ms raised-cosine ramp at each edge and ATTACK / RELEASE lengthen those ramps as a percentage of the grain's lifetime. When one grain hands over to the next, the successor starts as its predecessor enters release and the complementary cosine ramps sum to one - a true content crossfade rather than a butt splice (amplitude-complete joins that still splice unrelated waveform are the classic source of granular buzz). One further refinement comes straight from tape-editing practice: at the moment a grain's start position becomes final, the engine scans a ± 3 ms window for the nearest zero crossing in the source and starts the grain there, so every edge begins at near-zero amplitude. On percussive or steep material - a piano note-on under the read region - this removes the buzz of a transient's spectrum being re-fired at the grain rate; where no crossing exists locally (silence, DC) the original position stands and because identical requested positions snap to the identical crossing, the phase alignment of a zero-SPREAD freeze is preserved by construction.

ATTACK / RELEASE - the grain envelope's edge lengths, each as a percentage of the grain's lifetime, on top of the 10 ms minimum ramps.

PANO - per-grain stereo spread, constant-power.

P.RAND - per-grain random pitch detune.

OCT - half the grains shift up or down an octave.

5TH - half the grains shift up a fifth or down a fourth.

P.SHIFT / TAPE - selects how pitch is achieved and nothing else. TAPE re-pitches grains by resampling, like varispeed tape; P.SHIFT uses the live pitch-shifter so the read speed through the source is independent of the note. Neither mode moves the read position - POS remains the sole anchor.

The scrub display

While the oscillator is in GR mode its wave display shows the loaded waveform with a white POS playhead, a faint band marking the SPREAD region and the live grains drawn as they fire, so position, size and scatter are visible at a glance. The grain-view button switches to an all-voices view in which each held note's grains are drawn in that voice's own colour - the same hue the envelope visualiser uses for the voice - so chords read as distinct interleaved clouds. With TIMESTRETCH engaged the display adds bar and beat gridlines derived from the stretched tempo and a brief status band appears while a re-stretch is prepared.

The grain display

While an oscillator is in GR mode its wave display shows a grain scrub view: the sample waveform with a moving POS playhead marking the current read point and the active grains rendered as they fire, so you can see the density and spread of the cloud at a glance.

11.3.2.3 PARTIALS MODE (PA)

Partials mode resynthesises the loaded sample additively. On load, NeuralDream analyses the sample into a bank of harmonic partials - sine components tuned to the natural harmonic series of the detected root - and tracks how each partial's level changes through the sample. Playing a note rebuilds the sound from those partials at the pitch of the note, giving a smooth, alias-free, keyboard-tracked tone that behaves like a true additive oscillator rather than a sample being pitched up and down.

Because the sound is rebuilt from tuned partials, it is inherently clean and stable across the keyboard and it responds to the harmonic controls below in ways a raw sample never could. The trade-off is character: the model reconstructs the harmonic content faithfully but does not reproduce the broadband noise and air of the original recording, so Partials mode suits sustained, pitched, harmonic sources - vocals, strings, pads, brass, synths - better than noisy or percussive material.

TIMESTRETCH applies in Partials mode too - the partial analysis is rebuilt from the tempo-stretched audio, so resynthesis follows the project tempo just as grains do. The TAPE / P.SHIFT control and its algorithm buttons belong to grain playback and do not appear for a Partials oscillator.

The harmonic bands

The eight knobs across the panel are octave harmonic bands. Each governs the level of a band of harmonics grouped by octave above the fundamental - labelled by the lowest harmonic in each band: 1, 2, 4, 8, 16, 32, 64, 128. Band 1 is the fundamental and its nearest partials; each band above covers the next octave of the harmonic series. Every band runs from silence to twice unity, with 100% (the halfway position) leaving that band of harmonics at its analysed level - so the default, flat setting reproduces the sample's own spectrum and moving the knobs re-balances the tone octave by octave, from a hollow fundamental-only tone to a bright, harmonically rich one.

Three momentary buttons in the left column reset every band at once for the focused oscillator: R1 sets all bands to 100%, R.5 to 50% and R0 to 0% - quick starting points for building a spectrum up from silence or trimming one down from full.

Sweeps

Each oscillator also has two harmonic sweeps, SWEEP 1 and SWEEP 2. A sweep is a moving emphasis that travels across the harmonic bands: as you turn a sweep knob its focus glides from the low harmonics to the high ones, crossfading smoothly between adjacent bands rather than switching them abruptly - the same progressive-crossfade behaviour as the Octave Filter engine. With two independent sweeps you can set up two moving formant-like peaks in the spectrum.

Each sweep has a UNI / BI polarity toggle. UNI adds its emphasis in one direction; BI works symmetrically about the centre. Both the sweep positions and their polarities are ordinary modulation targets, so an LFO or envelope on a sweep produces a living, moving spectrum.

Position

POS scrubs through the analysis frames of the sample - the additive equivalent of the granular read position. At one setting it holds a single spectral moment of the sample; swept (for example by a clock-synced LFO) it scans the sample's evolving spectrum over time. Because the partials are retuned to the played note, POS changes the timbre without ever changing the pitch.

The partials display

In PA mode the wave display shows a dedicated partials view: the eight band levels drawn as translucent bars over a faint image of the source waveform, the two sweep markers as coloured lines - SWEEP 1 in blue, SWEEP 2 in amber, matching the colours of their knob labels - and a white POS playhead marking the current read point in the sample. The bars and markers move live with the knobs and with any modulation, so you can watch the spectrum evolve as it sounds.

The PROCESS Panel

When an oscillator is in Partial mode a PROCESS tab appears in the tab row beneath the wave display, alongside PARTIALS. It opens a panel of five processing controls that reshape the partial data as it plays - the additive counterpart to wavetable warping. Every control is per-oscillator, works in full stereo and has a value readout. Every control except ODD/EVEN is a standard modulation target, exactly like every other knob. ODD/EVEN can be moved only by the global strip, not by the synthesiser's own per-note sources (see 11.4.3). The processing is applied in real time to the analysed partials, so it costs almost nothing and layers freely with the band and sweep controls.

The top row holds four controls, with WIDTH centred beneath:

STRETCH (50–200%, default 100%) - scales the *spacing* of the harmonics while holding the fundamental fixed. At 100% the harmonics keep their natural whole-number spacing. Below 100% the partials are squeezed together for a denser, more vocal tone; above 100% they are spread apart into inharmonic, hollow, bell- and metallic-like timbres a plain harmonic oscillator cannot reach. This is the panel's most characterful control.

TILT (–12 to +12 dB/oct, default 0) - a spectral brightness slope pivoting on the fundamental. Positive values lift the upper partials for more air and edge; negative values roll them off for a darker, rounder tone. A quick, musical brightness control that leaves the fundamental untouched.

SHIFT (–500 to +500 Hz, default 0) - adds a fixed frequency offset in hertz to every partial. Unlike pitch, a linear shift breaks the harmonic relationship between partials, turning a harmonic tone progressively clangorous, bell-like and inharmonic as you move away from centre.

ODD/EVEN (–100 to +100, default 0) - crossfades emphasis between the odd and even harmonics. Toward –100 only the odd harmonics remain, for a hollow, clarinet-like colour; toward +100 only the even harmonics, for a fuller, rounder tone; at 0 both are present as analysed.

WIDTH (0–200%, default 100%) - the stereo width of the partials. At 0% the output is mono; 100% is the width as analysed from the source; above 100% the stereo image is widened.

AIR (0–200%, default 0) - the residual noise layer. When NeuralDream analyses a sample into partials it separates the sound into two parts: the tuned harmonic partials and the broadband residue left over - the breath, air, bow noise and grit that has no fixed pitch and so cannot be carried by the harmonic model. The partials rebuild the pitched tone; AIR brings that residue back. At 0 the residual is silent and costs nothing - the noise layer is not computed at all - at 100% it plays back at its analysed level and above that, up to 200%, it is exaggerated for extra breath and texture.

This is the control that answers the trade-off noted at the top of Partial mode: that the model reconstructs the harmonic content faithfully but does not, on its own, reproduce the broadband noise and air of the original recording. A little AIR restores the sense of a real instrument breathing behind the tone; more of it foregrounds the noise for an airy, hissing or gritty character.

The residual is not a generic hiss but a reconstruction of the sample's own noise. It is rebuilt from a bank of band-filtered noise matched to the noise spectrum captured in the analysis, it follows the analysis frames as POS scrubs so the air evolves with the recording and it responds to the eight harmonic band knobs, so shaping or sweeping a band shapes its air alongside its partials. It is reproduced in full stereo - a stereo source is given decorrelated left and right noise so the air keeps the width of the original, while a mono source shares one stream across both sides. Like every PROCESS control it is a modulation target, so the breath can swell under an envelope, LFO or expression - route Pressure to AIR, for instance and leaning into a held note pushes more air as you press.

AIR only has something to add on material that was analysed with a residual layer. Pitched, harmonic sources that carry breath or air - vocals, wind, bowed strings - gain the most; a clean synthetic tone with no noise in the first place gives it little to restore.

At their default settings (STRETCH and WIDTH at 100%, the rest at 0) the PROCESS panel is transparent - an untouched panel leaves the partials exactly as analysed.

11.3.2.4 MODAL MODE OSCILLATORS (MO)

Modal mode does not read anything. Where the other three oscillator modes play back material - a wavetable, a sample as grains, a sample as partials - Modal mode generates its sound by striking a resonator and letting it ring. It is the difference between playing a recording of a bell and hitting one.

The idea is simple enough to state in a sentence. Sound is sent into a delay line and part of the output of that delay line is fed back to its input. Whatever frequency exactly fits the round trip comes back in step with itself and grows louder; everything else comes back out of step and cancels. Make the round trip short enough and the delay line stops sounding like an echo and starts sounding like a physical object with a natural pitch. Feed it a short click and it sounds plucked. Feed it continuous noise and it sounds bowed or blown. Take the input away entirely and, if the loop is fed back hard enough, it carries on singing by itself. A Modal oscillator does this up to six times over, because one resonance on its own is only a comb filter and settles to a plain tone; a real bell, plate or drum rings on several at once and it is their interference that keeps a struck note alive. How many and how far apart they sit, are yours to set with MODES and SPREAD.

Each of those resonances is modelled on the analogue bucket-brigade delay, a chain of tiny capacitors that pass a charge along in step with a clock. Faster clock, shorter delay, higher pitch. That single fact carries a behaviour that a conventional filter-based resonator does not have and it is worth knowing because it shapes how the instrument plays across the keyboard:

A low note rings for a long time and a high note rings briefly.

Nothing in NeuralDream implements that. It is what happens. The loop loses a fixed proportion of its energy on every trip round; a high note completes its trip more often each second, so it loses that proportion more often. The result is an instrument that behaves like a real resonant body without being told to and it means a single patch played low sounds like a long sustaining string while the same patch played high sounds like a short percussive tap.

The two halves

Every Modal Mode oscillator has an EXCITER and a RESONATOR and it helps to think of them as the hand and the object. The exciter is what strikes; the resonator is what sings. They have their own control panels beneath the wave display and the six controls that shape the sound most while you are playing sit on the oscillator panel itself.

Between the two sits the INPUT FILTER, which is the boundary of the instrument. Everything the exciter produces passes through it before reaching the resonator and it also shapes what you hear on the dry path, so it colours both halves at once.

One difference from every other oscillator mode is worth flagging early. The resonator always follows the keyboard, because it is the thing defining the pitch. The exciter does not have to. That decoupling is deliberate and it is the source of most of the unusual sounds this mode can make - it is described under TRACK below.

The oscillator panel - the six main controls

The upper half of the oscillator panel is unchanged from every other mode: FREQUENCY and FINE, the octave selector and the F1 / F2 filter routing all behave exactly as they do in Wavetable, Granular and Partial mode. Below them sit the six Modal controls, laid out in the order the sound actually travels - across the top row first, then the bottom.

WAVE - the shape of the exciter, morphing continuously from a pure sine, through triangle and sawtooth to a narrow pulse and finally to a square. It decides how much harmonic material you are throwing at the resonator. A sine barely wakes the resonator up and gives a soft, hollow result. The narrow pulse at around three quarters of the way up is the sharpest thing the exciter can produce and hits the resonator with everything, so that is where it rings hardest and brightest; the square at the top is fuller and rounder than the pulse rather than sharper still. Because the morph is continuous rather than a set of choices, it is a modulation target like any other knob, so the character of the strike can move under an envelope or LFO.

INV - the feedback polarity and it is a switch rather than a knob because it does one thing decisively. An inverted loop reinforces a different set of frequencies from a normal one: the resonance settles an octave below where it sat before and only the odd harmonics of that lower note survive. Round and direct becomes hollow and nasal - the difference between a plucked string and a stopped organ pipe. Expect the pitch to drop and the level to fall when you switch it in, because the exciter is tuned to the note you played and the inverted loop cancels at exactly that frequency; raise FB or bring the exciter down with EXC FREQ to work with it rather than against it. It is a single change that alters the whole personality of the resonator and it is worth trying on any patch that feels too plain.

IN FILTER - a gentle low-pass filter between the exciter and the resonator. It decides how bright the strike is and therefore how the resonator responds to it: brighter in gives a sharp, metallic, bell-like resonance, darker in gives something round and wooden. At the very top of its travel it opens completely and passes the exciter through untouched, so fully up genuinely means out of the way rather than nearly out of the way. Because the dry path passes through it too, it also shapes what you hear when DRY/WET is turned down. If a patch sounds harsh, this is usually the control to reach for before the filters.

FB - feedback - how much of the resonator's output returns to its input and the single most powerful control in the mode. Low settings add a little body and colour to the exciter. Higher settings make the resonator sing properly and hold the note after the exciter has finished. At the very top it produces more energy than it loses and continues sounding on its own with no exciter at all, which is a sound world of its own - see Self-oscillation below.

FBF - the feedback filter - a gentle low-pass inside the loop. Because it acts on every single pass rather than once, it dominates how the sound decays: it sets how long the note rings and how much brightness is lost each time round. It does not affect the pitch. Earlier versions of the plugin went slightly flat as this control was closed, because the delay the filter itself contributes was not being accounted for in the loop's length; that is corrected and the resonator now lands on the played note to within a tenth of a cent wherever FBF and DAMP are set.

WET - the dry/wet balance - what you struck against what rang. Fully dry you hear only the exciter, filtered by IN FILTER; fully wet you hear only the resonator. It is also how you play a self-oscillating patch: with the exciter silent the resonator is already sounding, so this control becomes its volume.

11.3.2.4.1 THE MODAL EXCITER PANEL

Select MO on an oscillator and the tab row beneath the wave display changes to MIXER, EXCITER and RES. The WARP, GRIT, GRAINS, PARTIALS and PROCESS tabs are hidden, because none of them apply to a mode with no sample and no wavetable. As with every other mode, each oscillator remembers which tab it last had open. The EXCITER panel holds everything about the thing doing the striking, in two rows: what makes the sound on top and how it is shaped underneath.

Row one - the sound source

SC/O3 IN - external input. This crossfades an outside signal against the internal exciter, so that whatever arrives becomes the hammer that strikes the resonator. At zero it is completely absent and the oscillator behaves normally; at full it is the entire exciter. The SC SRC button below chooses the source: the DAW's sidechain or OSC 3.

It passes through the exciter envelope like any other source and that is the important part: the input is gated by the notes you play, so holding a chord tunes several resonators at once and lets the incoming audio strike all of them together. Send a drum loop in and the kit plays your harmony.

Routing OSC 3 in is the more far-reaching of the two. It puts a full wavetable, grain cloud or partial bank in front of the resonator - an exciter with its own filter, envelope and modulation, rather than the waveform-and-noise pair the panel offers. OSC 3 follows the keyboard like any oscillator, so its partials land on the resonances rather than between them and its waveform audibly changes the result: a narrow pulse gives a bright, hollow strike where a sine gives a

round one. Pull OSC 3's mixer fader to silence and it still drives the resonator, so it can be an exciter and nothing else.

SC SRC - what SC/O3 IN listens to: the DAW sidechain or OSC 3. On OSC 3's own panel this is fixed at Sidechain, since an oscillator cannot feed itself.

HAMMER / PLUCK - which of the two exciters is in use. See section 11.3.2.4.1.1 below.

EXC FREQ - the exciter oscillator's own pitch, spanning eleven octaves centred on the note it is striking. This is the pitch the exciter plays in its own right, before the keyboard is allowed to move it. At the halfway point, where it sits on a fresh patch and with TRACK fully up, it plays exactly the note you play, which is why a new patch behaves like an ordinary instrument. Because the note sits at the centre of the travel rather than off to one side, a bipolar modulation route has the same range to work with in each direction - though with eleven octaves on the control that range is wide, so modulation depths that felt right on a narrower range will now be large.

Wind TRACK down and this becomes the fixed pitch the exciter sits at while the resonator climbs with the keys; with TRACK at zero the travel runs from 20 hertz to 12 kilohertz. The lowest tenth or so of the travel sits at the 20 hertz floor, which is the price of keeping the played note exactly at the centre.

The exciter's waveforms are band-limited, which is what allows the range to reach 12 kHz at all. Without it, a ramp or pulse played that high would fold back as inharmonic partials - the fizzing, gritty quality that gives away a naive oscillator - rather than simply sounding high.

EXC FINE - fine tuning for the exciter alone, three semitones either way. It is summed with EXC FREQ. Detuning the exciter slightly against the resonator is the most direct route to the slow beating this mode is good at, because the two are then feeding and ringing at very slightly different pitches. A few cents is usually enough.

NOISE - blends the exciter away from its waveform and towards a noise source. Pure waveform gives a clean, pitched strike; pure noise gives a breath or a bow. This is how you move from plucked and struck sounds towards wind and bowed-string behaviours and a small amount of noise mixed under the waveform adds air to an otherwise pure strike.

S&H RATE - shapes that noise by sampling and holding it. Fully up it is smooth white noise. Wind it down and the noise is held in steps, so it breaks up into a gritty, stuttering, almost pitched rattle and right at the bottom the steps are slow enough to hear individually as a low clatter. The part worth knowing is what governs the step rate over most of the travel: for the first three quarters of the knob the clock is tied to the exciter's own pitch and spans eight octaves around it, four below and four above. That means the rattle tracks the keyboard whenever the exciter does, so stepped noise stays musical as you play up and down instead of sitting at one fixed rate. Only the last quarter of the travel leaves the exciter behind and climbs steeply into continuous white noise. Because the clock follows the exciter, EXC FREQ and TRACK change the character of the noise as well as the character of the waveform.

DRIVE - how hard the exciter hits the resonator. Gentle settings give a polite, well-behaved resonance. High settings push the delay line into its own saturation, which thickens and dirties the tone in the way that overdriving real hardware does. It is a level control and a character control at the same time.

TRACK - exciter tracking and the control with no equivalent anywhere else in the synthesiser. The resonator always follows the keyboard. The exciter only follows it as far as this control asks. At 100% the two move together and the sound is coherent and instrument-like across the whole keyboard. Wind it down and the exciter falls back towards its own fixed pitch, the one you set with EXC FREQ and EXC FINE, while the resonator continues to climb with the keys, so the harmonics being fed in slide across the resonator's response as you play. At zero the exciter sits on that pitch whatever you play, which is middle C until you move EXC FREQ. The result is beating, shifting spectra and interference patterns that change character note by note - impossible to obtain any other way in NeuralDream and the first thing to try when a Modal patch feels too well behaved.

Row two - the exciter envelope

TRACK opens this row and the five controls after it are the exciter envelope. They shape the exciter only, not the resonance. This is the point most worth understanding about the mode: when the envelope finishes, the exciter stops but the resonator carries on ringing exactly as a struck object does when you take your hand away. How long it rings is set by FB and FBF, not by these controls.

ONSET - silence between the key going down and the exciter starting, up to two seconds. On one oscillator it sounds like a late entry. Set differently across the three oscillators it gives strums, flams and rolls from a single key press. A note released during the onset never sounds at all.

ATTACK - how quickly the exciter reaches full level. Short is a strike or a pluck; long is a bow or a breath being taken.

DECAY - how quickly the exciter falls from full level to its sustain level.

SUSTAIN - the level the exciter holds at while the key is held. At zero the exciter is a single hit and everything you hear afterwards is the resonator ringing - the classic plucked or struck sound. Raised, the exciter keeps feeding the resonator for as long as you hold the note, which is how bowed and blown sounds are made.

RELEASE - how quickly the exciter fades after the key is released. The resonator continues regardless.

Beneath the two rows sit four switches:

ENV RANGE - the timescale the four envelope knobs cover - SHORT for percussive work, MEDIUM for most sounds and LONG for slow swells. Set the range first and then dial the envelope, rather than the other way round.

HOLD - the exciter ignores note-off and keeps feeding the resonator until you take the note away by other means. Intended for drones and beds.

LOOP - the exciter envelope repeats for as long as the note is held, running its onset, attack and decay again and again. The first pass rises from silence and every later pass starts from the sustain level, so the repeats are pulses riding on a held tone rather than full re-strikes. This turns the exciter into a rhythmic pulse generator - tremolo, buzz rolls and repeating strikes - with no sequencer involved. Combined with ONSET set differently across the three oscillators it is where most of the rhythmic Modal patches come from.

LEGATO - changes what happens when notes overlap. With LEGATO off, every note starts its envelope from silence. With it on, a note arriving while the previous one is still in its release picks the envelope up from wherever it had got to, rather than dropping to zero and starting again. The result is smooth, connected phrasing instead of a click at every note change and it is what you want for pads and legato lines. Switching LOOP on enables LEGATO automatically, because a repeating envelope only makes sense as a continuous one.

11.3.2.4.1.1 HAMMER AND PLUCK

The EXCITER panel offers two ways of setting the resonator going and they are different in kind rather than degree. HAMMER is the original and remains the default, so every existing preset is untouched.

A **HAMMER** is a force applied over time. The exciter drives the resonator continuously for as long as the envelope holds it open, which is why a hammer patch with any sustain rings indefinitely - it is being refuelled. Everything on the panel shapes that force: the waveform, the noise, the sample-and-hold, the envelope.

A **PLUCK** is not a force at all. It is an initial condition: the string is displaced into a shape and released and everything after that is the resonator losing what it was given. Nothing drives it. This has consequences that are worth stating plainly, because they are not what a synthesiser player expects.

The envelope does nothing. There is no drive for it to shape, so ATTACK, DECAY, SUSTAIN and RELEASE grey out, along with WAVE, NOISE, S&H RATE, EXC FREQ, EXC FINE and TRACK - all

of which shape an exciter oscillator that a pluck does not have. FEEDBACK greys out too, replaced by DECAY.

PICK becomes the tone control. It sets where along the string the displacement is applied and that decides which resonances receive energy and which do not - plucking at the midpoint silences the even ones, exactly as it does on a real string. Crucially it does this at constant loudness: PICK changes the tone without changing the level, which no other control in this mode does.

DRIVE becomes pure loudness - how far the string is pulled before release. On a linear string, pulling harder makes it louder, not brighter.

DECAY is set in seconds rather than as a percentage. It is pitch-compensated, so a one-second decay is one second at every note on the keyboard; a feedback percentage cannot do this, because a low note makes far fewer round trips per second than a high one. It governs the whole note, including any external input arriving through SC/O3 IN.

ONSET still delays the strike and delays it for both the pluck and anything arriving on SC/O3 IN.

The practical difference is dramatic at the attack. A pluck has its energy in the string before the first sample is computed, where a hammer has to fill the delay line one round trip at a time - more than a thousand times the energy in the first millisecond. That is what a pluck sounds like and no envelope setting on a hammer will produce it.

Note also that a pluck needs far more DECAY than a hammer needs FEEDBACK to ring for the same length, precisely because nothing is refuelling it. A hammer patch converted to PLUCK will sound short until DECAY is raised.

11.3.2.4.2 THE MODAL RESONATOR PANEL

The RES tab holds the body itself - how many resonances it has and how they are spaced - together with its tuning and the analogue character of its delay lines, in two rows.

Row one - the body and its tuning

MODES - how many resonances ring at once, from one to six. This one is a button rather than a knob and it sits beneath the row of knobs; each click steps it on by one and it reads back what it is set to, so it shows 4 MODES or 1 MODE when there is only one. At one the oscillator is a single resonator and behaves like a plain plucked or bowed string: clean, focused and settling quickly to a simple tone. Raising it adds further resonances and the sound gains body, complexity and the slow beating between partials that makes a struck object sound like an object rather than a note. Six is a full struck body. It costs a little more processing per voice, which is why it is adjustable rather than fixed.

SPREAD - where those resonances sit and the control that decides what kind of object you are hitting. At zero they are spaced in whole-number ratios - the harmonic series - and the bank reads as one thick, unified string or tube. Wind it up and they move apart towards the ratios of a struck bar, which are not whole numbers at all and the sound turns to bell, plate and struck metal. The pitch stays where you played it throughout, because the lowest resonance always leads; what changes is everything around it. The readout tells you directly how much stretch you have dialled in: 0.0 semitones at rest, where the resonances sit in whole-number ratios, rising to 12.0 semitones at the top. The absolute interval from the note to the second resonance is an octave plus whatever this shows. This is the fastest route to bell tones and it is worth sweeping slowly with a chord held down to hear the object change shape.

OFFSET - transposes the resonator in exact semitones, up to an octave either way, without moving the exciter. This is how you tune the resonant body to an interval against whatever is striking it - a fifth, an octave, a seventh - which changes the colour of the result rather than its pitch.

FINE - the same idea continuously, within three semitones either way. Use it to detune the resonator against the other oscillators for width or against another Modal oscillator so the two bodies beat slowly against each other.

GLIDE - how long the resonator takes to slide from one note to the next, up to one second. Because only the resonator glides and not the exciter, the two separate briefly during the slide, which is a pleasing effect in itself.

TRACK - resonator tracking. At 100% the resonator follows the keyboard exactly and the mode plays as a tuned instrument - a string, a tube, a bell. Wind it down and the resonant body falls back towards a fixed pitch of middle C while you play up and down the keys, so the pitch barely moves and what changes instead is which part of the body you are striking. This is how drums, plates and struck metal are made: one object, hit in different places, rather than one note transposed. At zero the body sits on middle C whatever you play and the keyboard becomes a way of choosing the strike rather than the note.

Row two - the bucket-brigade character

These four controls are what turns a clean digital delay line into something that behaves like a physical object in a real circuit. The first two govern how the body responds: how quickly its bright harmonics fade relative to its low ones and where along it the strike lands. The second two are the analogue faults proper, the noise and the instability of the hardware. All four rest at zero, where the resonator is precise, predictable and completely repeatable and all four are modulation targets, so the character of a patch can itself move over time.

DAMP - how quickly the high harmonics die away compared with the low ones. FB FILTER sets the overall colour of the loop and DAMP adds damping on top of it, so the two work together rather than one scaling the other: whatever DAMP is set to, FB FILTER still governs the colour and whatever FB FILTER is set to, DAMP still governs how fast the top of the sound leaves it. It rests at zero, where it adds nothing and the loop is FB FILTER's alone. Turning it up makes the high harmonics die progressively faster than the low ones, so the note starts as bright as it ever did and dulls as it rings, which is what a struck wooden or damped object does. Across the whole knob the energy above 1.2 kilohertz moves by about fourteen decibels against the energy below it, while the note's own level barely changes, so this shapes tone over the life of the note rather than loudness. If you want the sound brighter than DAMP at zero gives you, open FB FILTER: that is the control that sets the starting colour.

PICK - where along the body the strike lands and the control that changes the voicing of the mode most for the least effort. Strike a string at a point one third of the way along and you cannot set the third harmonic moving, because that is exactly where that harmonic stands still; strike it halfway along and every even harmonic goes the same way. This control does the same thing to the resonator, sweeping a notch through the harmonics as you turn it. At zero the notch sits above the fiftieth harmonic, so it is out of the way and the strike is hard and bright, as though you had hit the very end. Turning it up brings the notch down through the spectrum; at the top it sits on the second harmonic and takes out every even one, which is soft and hollow, the sound of striking the middle. The level does not move as you turn it - the fundamental holds within a tenth of a decibel from end to end - so what you hear is purely the change of voicing. Sweep it slowly on a held note and you can hear individual harmonics drop out and come back.

CLOCK - the noise the clocking drags along with it. A bucket-brigade line is a noisy device by nature and longer delays are noisier still. The hiss here follows the resonator: it rises as the body rings and falls away with it, so it sits under the note as part of the sound rather than running underneath the whole patch as a constant background. That gentle, moving hiss is a large part of why these units are described as warm rather than clean. Raise this control for air, grit and vintage character. It is worth knowing what this control deliberately does not do: the clock itself is far above hearing at any note you will play and every well-built delay filters away what little of it strays into the audio path, so there is no whistle here and none should be expected.

DRIFT - nothing analogue holds still. The clock wanders slowly, so the pitch that the loop settles on wanders with it, producing gentle detuning, beating and movement that never quite repeats. This is the control that stops a held Modal chord sounding static and a small amount is worth having on almost any patch.

11.3.2.4.3 SELF-OSCILLATION

With FB near the top the resonator produces more energy than it loses on each pass. At that point it no longer needs to be struck at all: the circuit's own noise floor is enough to start it and the loop amplifies that into a sustained tone which continues until the feedback is reduced or the note ends.

This gives Modal mode a second way of being played. Turn DRIVE down so the exciter is silent, raise FB to the top and use WET as the volume control for the resonator's own tone. Because WET is a modulation target, an envelope routed to it lets you bow or strike the self-oscillating resonator directly and the pitch is still set by the keyboard, NOTE OFFSET, FINE and GLIDE. The result is a pure, singing, slightly unstable tone that is quite unlike anything the other three oscillator modes produce.

11.3.2.4.4 GETTING STARTED WITH MODAL

Five starting points, each of which is a few controls away from the others:

A plucked string - WAVE around a third of the way up, NOISE at zero, ENV RANGE on SHORT, ATTACK at minimum, DECAY short, SUSTAIN at zero. FB high, FBF around two thirds, WET fully wet. Play low for long sustains and high for short taps - the ring time follows the keyboard on its own.

A bowed sound - From the plucked patch, raise NOISE towards the top, lengthen ATTACK, raise SUSTAIN so the exciter keeps feeding while the note is held and set ENV RANGE to MEDIUM. Lower IN FILTER a little to take the edge off.

A struck drum or plate - From the plucked patch, wind the resonator's TRACK down towards zero so the body stays at a fixed pitch however high or low you play, raise MODES to five or six and set SPREAD around half so the partials are inharmonic without being bell-like. Shorten DECAY further and raise DRIVE for a harder hit. Take DAMP up towards three quarters so the brightness falls away fast and leaves a thud rather than a ring and set PICK around a third of the way up to knock the higher partials out of the strike.

A bell - Set MODES to six and SPREAD high - three quarters or beyond - so the resonances take up the ratios of a struck bar. Keep the strike short: ATTACK at minimum, DECAY brief, SUSTAIN at zero. Raise FB well up so the body rings on after the strike and open FBF so the upper partials survive long enough to shimmer. A little DRIFT sets the partials beating gently against one another, which is what stops a long bell sounding static.

A sidechain-driven resonator bank - Route a drum loop to the plugin's sidechain, raise SC IN to full, set SUSTAIN high so the exciter gate stays open while you hold notes and play chords. Each note you hold is a resonator tuned to that pitch and the loop strikes all of them at once.

Two final notes: First, Modal is a full oscillator mode, so all three oscillators can run it at once - detuned against each other with FINE or set to different NOTE OFFSETs, they build a genuine multi-mode resonant body in the way that a piano string is really three wires. Second, every Modal control is an ordinary oscillator control: a modulation target with a value readout, saved and recalled with the preset and available to the DAW for automation exactly like every other knob in NeuralDream. The one exception is NOTE OFFSET, which is a setting in whole semitones rather than something to sweep and can be moved only by the global strip (see 11.4.3).

11.3.2.4.5 FM AND SYNC IN MODAL

The synth's FM and SYNC controls work in Modal mode and they act on the exciter - the one part of the mode that has an oscillator phase to work with. They are not available for Granular or Partial, which have no phase at all and the controls grey out accordingly.

FM modulates the shape of the exciter rather than its pitch, so what it changes is how much harmonic material is being thrown at the resonator. Small amounts add bite and edge to the strike; larger amounts turn a simple waveform into something clangorous, which the bank then rings on. Think of it as a timbre control on the hammer.

SYNC restarts the exciter mid-cycle on every cycle of OSC 1, exactly as hard sync does on a wavetable oscillator. The resonator hears a discontinuity at each restart, which is bright, edgy and aggressively harmonic - and because the resonator then rings on it, the result is more sustained and less brittle than hard sync usually sounds.

11.3.2.5 IMPORTING AUDIO

Any audio file - a vocal phrase, a string note, a synth loop, a field recording - can become a scannable wavetable. Drop the file onto an oscillator's wave display and choose WAVETABLE from the dialog; NeuralDream copies it into your user wavetable library under an Imported category (transcoding to WAV where needed, so the library's scanner can index it), converts it and points the oscillator at the result. From that moment it behaves exactly like any factory table: Position sweeps it, the 3D view shows its terrain, warping and modulation apply and it appears in the wavetable browser for every preset thereafter.

Trimming silence on import

Many recordings arrive wearing silence - a breath before the downbeat, a long reverb tail after the final note. The conversion engines are faithful to what they are given: silence in the file becomes silence in the result and a recording that is one-third quiet spends one-third of its Position range on flat, empty frames. The two switches on the import dialog, TRIM START and TRIM TAIL, remove that silence before conversion. Both are off by default, they operate independently and they apply to all three import routes - grains, partials and wavetable alike.

The threshold is musical rather than absolute: material more than 60 dB below the file's own peak is treated as silence, so a quiet recording is judged against itself, never against full scale. Detection measures short RMS windows from each end of the file inwards, trimming only at the ends and never within - a pause in the middle of a phrase is always preserved - and each cut point is nudged to the nearest zero crossing with a small safety margin, so nothing clicks and no transient is shaved. A file containing no qualifying silence is left untouched.

The switches are separate for a musical reason. A drum loop with a soft count-in wants TRIM START, so beat one lands at the very start of the result - but its tail is part of its length and its length is its tempo: trim the tail of a loop and it no longer spans a whole number of beats. Trim the start, keep the tail. A sustained note with a long reverb decay is the opposite case: TRIM TAIL reclaims the Position range the tail was wasting, so the whole sweep is live timbre.

Every import reports its outcome in a brief on-screen notice: the exact amounts removed and the resulting length ("Trimmed tail 2.8s - 8.00s → 5.2s"), a note that no silence was found or an explicit message if the trim could not be performed. An import never trims silently and never fails silently.

Finally, re-imports replace rather than accumulate. Importing a file whose name matches an existing entry in your Imported library replaces that entry - there are no numbered duplicates - so you can audition the same recording trimmed and untrimmed or re-import after editing the file externally and the oscillator always plays the version you imported last.

The conversion begins with classification, because a file that is already a wavetable must not be re-analysed. Three tests identify a genuine single-cycle table: a total length that is an exact multiple of the 2048-sample frame; measured cycle lengths close to one cycle per frame; or a Serum-style 'clm' metadata chunk declaring the frame size, which is honoured directly - Serum-format tables import frame-for-frame. Anything that fails these tests is treated as an ordinary recording and the classifier is deliberately biased in that direction: a wavetable mistakenly resynthesised still plays correctly, whereas a recording mistakenly read as raw frames produces the broken, aliased "grass terrain" familiar from naive importers.

A recording is then rebuilt by harmonic resynthesis, in the tradition of sinusoidal analysis-synthesis modelling (McAulay and Quatieri, "Speech Analysis/Synthesis Based on a Sinusoidal Representation," IEEE Transactions on Acoustics, Speech and Signal Processing, 1986). The audio is examined in overlapping analysis windows stepped through the file; in each window the harmonic content - the amplitudes and phases of the partials above the detected fundamental - is measured and from those measurements one perfect 2048-sample single cycle is synthesised.

Each window therefore contributes one frame and the frames in order replay the timbral journey of the recording. The analysis hop adapts to the material: up to 1024 frames are produced, with the hop never finer than 6 ms, so a short one-shot yields a modest stack of frames while a long evolving source spreads across the full thousand - and in every case the frames cover the whole file, so Position at 0 is the recording's opening timbre and Position at 100% its final one.

Two details are worth knowing because they are audible. Stereo is preserved: the left and right channels are analysed and synthesised independently but each frame receives a single shared phase-alignment rotation, which keeps the stereo image coherent and prevents adjacent frames from fighting each other as Position scans - a swept import stays smooth rather than phasey. And the root pitch is measured from the audio itself using the harmonic product spectrum (Noll, "Pitch Determination of Human Speech by the Harmonic Product Spectrum, the Harmonic Sum Spectrum and a Maximum Likelihood Estimate," 1970), with any note name in the filename used only as a corroborating witness - so an imported cello D plays a D when you press D, regardless of what the file was called.

The musical result: a held key with static Position is a frozen, perfectly pitched single cycle from one instant of the recording; a modulated Position is the recording's timbre replayed at any speed, forwards, backwards or scrubbed by an LFO - pitch and time fully decoupled, the classic wavetable trick applied to any sound you can drop. Imports pair especially well with the warping engines and an imported table also serves as a Fold Type source (see Wavetable Warping), so one recording can both play and shape.

11.3.3 THE PANEL TABS

Beneath the wave display is a row of tab buttons. These switch the area below the display between the different control panels for the focused oscillator, so a single space serves the mixer, the warping engines and the sample-mode panels in turn. Only one panel is shown at a time and pressing a tab brings its panel to the front.

Switching an oscillator's mode swaps the mode-specific tabs to match - so moving from Wavetable to Partial replaces WARP and GRIT with PARTIALS and PROCESS and moving to Modal replaces them with EXCITER and RES. If the panel you were viewing belongs to a mode you have just left, the view falls back to MIXER so you are never left looking at a panel that no longer applies. Each oscillator also remembers which of its tabs you last had open, so returning focus to it restores that panel.

Two further view toggles share the same row - 2D/3D and OSC/SUM - but these control the *wave display itself* rather than opening a panel and they appear only for Wavetable oscillators. They are described with the wave display in Section 8.x.

Which tabs appear depends on the mode of the focused oscillator (set by its WT | GR | PA | MO toggle), because each set of panels only applies to one mode. The row always begins with MIXER, then shows the tabs relevant to the current mode:

MIXER - always present. The four-source mixer with per-channel level and pan (Section 11.3.4). This is the default panel.

WARP and GRIT - shown for a Wavetable oscillator. They open the two warping engines: WARP for PD and Fold, GRIT for saturation and lo-fi (Section 11.3.2).

GRAINS - shown for a Granular oscillator. It opens the granular controls: DENSITY, SIZE, SPREAD, POS and the FILL toggle (Section 11.3.2.2).

PARTIALS - shown for a Partial oscillator. It opens the harmonic-band panel: the eight octave band levels, the R1 / R.5 / R0 resets and the two sweeps with their UNI / BI polarity toggles (Section 11.3.2.3).

EXCITER and RES - shown for a Modal oscillator. EXCITER opens the sidechain, the exciter's own pitch and fine tuning, the noise and drive controls and the exciter envelope with its TRACK control and its ENV RANGE, HOLD, LOOP and LEGATO switches; RES opens the modal bank with its MODES button and SPREAD, the resonator tuning and the four character controls that

give the body its damping, its strike position and its analogue noise and drift (Section 11.3.2.4). A Modal oscillator shows no other mode tabs, because it reads neither a sample nor a wavetable.

PROCESS - shown for a Partial oscillator alongside PARTIALS and now for a Granular oscillator alongside GRAINS. On a Partial oscillator it opens the full post-partial processing panel; on a Granular oscillator it opens a reduced version holding the shared sample-trim controls together with the TIMESTRETCH and TAPE / P.SHIFT playback options (11.3.2e). A PROCESS tab appears whenever an oscillator is in Partial or Granular mode. In Partial mode it opens the post-partial processing controls that reshape the partial data as it plays; in Granular mode it opens the sample-region and playback controls. Every control is per-oscillator, works in full stereo and has a value readout. Every control except ODD/EVEN is a standard modulation target, exactly like every other knob. ODD/EVEN can be moved only by the global strip, not by the synthesiser's own per-note sources (see 11.4.3).

PRESETS - always present and always the last button in the row, to the right of the keyboard icon. It opens the synthesiser's own preset list in the area below the display. It used to sit in the preset strip. That strip holds the controls for the whole plugin while these presets belong to the synthesiser, which is why PRESETS now sits with the synthesiser's other tabs. The buttons in this row are spaced evenly across the width of the display.

OFFSET (0-100%) and **INSET** (0-100%) trim the usable region of the loaded sample: **OFFSET** moves the start point inward from the first sample, **INSET** moves the end point inward from the last. Between them they define the window of the recording the oscillator reads - a quick way to crop silence from the head or tail, isolate a phrase or tighten a loop. They are shared by both Granular and Partial mode, since they act on the sample region itself rather than on the mode and both are modulation targets, so the read window can move under an envelope, LFO or expression.

TIMESTRETCH - tempo matching. A TIMESTRETCH toggle with an editable SOURCE BPM field sits at the top of the PROCESS panel, in both Granular and Partial mode. With it off the sample plays at its native speed. Turn it on and the sample is stretched to the host tempo, so a loop recorded at one tempo locks to your project's tempo without changing pitch. The SOURCE BPM field holds the sample's own tempo - NeuralDream seeds it from the filename on import where it can and you can correct it by typing - and the amount of stretch follows the ratio between that and the host tempo. The stretch always works from a pristine copy of the original recording, so toggling it or changing the source tempo never accumulates artefacts. The stretch follows the project. Change the host tempo and the sample re-stretches to match automatically - no need to touch the button - settling about half a second after the tempo comes to rest. During a continuous tempo ramp (a *rallentando* or *accelerando*) the sample holds at its last settled tempo and snaps to the final tempo once the ramp ends, rather than chasing every intermediate value. This tracking works whether or not the plugin window is open, so a saved session opens already in time and because presets store the untouched recording plus its source tempo - never a stretched copy - a preset made at one project tempo locks itself to whatever tempo the next project runs at. While a stretch is pending the wave display says TIME STRETCHING and if the source tempo is unknown it says SET SRC BPM TO STRETCH; when neither message shows, what you hear is exactly what the panel states.

TAPE / P.SHIFT - grain playback mode. Beneath TIMESTRETCH a TAPE / P.SHIFT toggle sets how grains are pitched. In TAPE mode - the classic granular behaviour - pitch and read speed are linked: a higher note reads through the sample faster, exactly like speeding up tape, so the scan through the recording quickens as you play up the keyboard. In P.SHIFT mode the two are decoupled: every note reads the sample at the same speed and the pitch is applied by a true time-domain pitch shifter, so the recording scans at one constant rate whatever note you play while still tracking the keyboard. TAPE suits classic granular character and pitch-following scans;

P.SHIFT suits sustained or tonal material where you want the sample's own movement to play back at a steady speed across the whole keyboard. P.SHIFT suits sustained or tonal material where you want the sample's own movement to play back at a steady speed across the whole keyboard. The pitch shift happens live inside each voice, so P.SHIFT engages instantly - there is nothing to prepare and a tempo re-stretch is heard immediately at every pitch. P.SHIFT algorithms - Beside the P.SHIFT button, two buttons choose how the live pitch shift is performed. DRUM is a dual-tap harmoniser: the lightest on the CPU and the default, at its best on percussive material, with a slight flutter on hard transients at large shifts. TONE is a phase vocoder with identity phase locking: smoother and more focused on sustained and tonal material, at a noticeably higher CPU cost, with transients softening slightly. The choice is per oscillator, is saved with presets and takes effect from the next note played. Hover either button for a summary tooltip.

11.3.4 THE MIXER

Four sources feed the mixer: the three oscillators and a built-in white-noise generator. The mixer comes before the filters, not after them, which is why its heading reads MIXER -> FILTERS. Each channel's Level therefore sets how much of that source goes into its filter and so how hard that source drives the filter. Each channel has its own Level and Pan, with a vertical signal meter beside each fader so you can see how much each source contributes at a glance. The meters read each source's peak level after its fader, with an instant attack and a smooth fall. They are a guide to each source's share of the mix; whether a filter is actually being overdriven is shown by the red overload lamp on that filter, described in Section 11.3.5.

The OUTPUT fader, coloured yellow, is the one control on the mixer panel that comes after the filters. It sets the synthesiser's final level, after the filters and the amplifier and it does not change how hard anything drives the filters.

Beneath each channel sit two buttons, S and M - the industry solo/mute convention. Solo (yellow) is additive across sources: solo two oscillators and both sound. Mute (red) silences a channel. When solo and mute collide, solo wins for audibility - a soloed source sounds even with its mute lit and the mute is remembered, reasserting itself the moment the last solo clears. This is the majority DAW behaviour, chosen so that soloing is always a one-click way to hear something. Both buttons can be lit at once; each channel's state is visible from its colours alone. A source that is muted (and not soloed) shows a sleep indicator across its caption.

Panning uses a constant-power law, so moving a source across the stereo field keeps its perceived loudness steady rather than dipping in the centre. The noise channel is useful for adding breath to pads, chuff to plucks or as a modulation source in its own right when filtered.

On a fresh INIT, OSC 1 (a sine) feeds Filter 1 and OSC 2 (a saw) feeds Filter 2, each at a Level of 50%, with both filters fully open and both filters' Trim at -5 dB. At those levels neither filter is overdriven, so a key pressed on INIT gives a clean two-oscillator voice and both red overload lamps stay dark. Every new synthesiser starts in exactly that state, not only the one you get from INIT: a brand-new instance of the plugin and any layer you set to Synth, both begin there.

11.3.5 THE FILTERS

The mixed signal passes into two independent filters, Filter 1 and Filter 2. Each filter has the same controls:

Model - LADDER, ZDF, COMB, LPG, SHIFT or VOWEL: the circuit itself. These are six genuinely different pieces of mathematics, described in full in section 11.3.5.1 below. LADDER is the default. Since version 2.15.69 it is a transistor ladder that saturates on its own when it is fed a hot signal, even with Drive at zero, so a preset made before then can sound grittier than it did when it was made. The red overload lamp above Trim shows when this is happening and Trim turns the level into the filter down to cure it. The choice of model is independent per filter, so Filter 1 can be a LADDER while Filter 2 is a VOWEL; the model may be switched freely, even while notes sound.

Type - which output is taken from the circuit: LP, BP, HP or MULTI. This applies to LADDER and ZDF only and greys out for COMB, LPG, SHIFT and VOWEL, which have no taps to choose between. MULTI is a LADDER feature and is unavailable under ZDF; selecting it there returns the control to LP.

Drive - LADDER only and hidden for every other model. It pushes the signal harder into the ladder's input stage, which saturates and adds harmonics a linear filter cannot produce. Turning it up raises the gain into that stage by up to eight times while the level is brought back down after it, so a sound does not simply get louder: quiet notes come up, loud ones are compressed and the tone thickens, as an overdrive behaves. The saturation in the resonance path is gentler and does not change with the control, so resonance and self-oscillation stay stable at any setting. Because the ladder saturates on its own, a hot signal is shaped by it even with Drive at zero; the red overload lamp shows how hard the ladder is being driven.

Trim - LADDER only and hidden for every other model. It sets the level going into the filter, from -12 dB to +6 dB, before Drive. Turn it down to take a hot sound out of overload or up to push a quiet sound into the ladder's grit. It raises or lowers the filter's output by the same amount, so bring the oscillator Level or OUTPUT back up afterwards if needed. At its default of 0 dB it has no effect, so every preset saved before Trim existed sounds as it did; INIT sets both filters' Trim to -5 dB. Trim sits at the outer end of the knob row, to the left of Drive in Filter 1 and to the right of Drive in Filter 2, so the two filters mirror each other.

Overload lamp - the red lamp directly above Trim, shown with the LADDER model. It glows in proportion to how hard the ladder is being driven: it stays dark while the filter is clean, begins to glow at the first audible saturation and is fully lit at heavy distortion, fading at the same rate as the mixer meters. It measures what actually reaches the ladder, after Trim and Drive and including the resonance, so turning Drive up lights it too, as it should. With several notes playing, it follows the hardest-driven one.

Bass - LADDER only and hidden for every other model, the button beneath LP Poles. It is shown only while the filter is a low-pass, which means TYPE set to LPF or TYPE set to MULTI with HP Poles at 0 or LP Poles at 4. Every 24 dB per octave resonant filter of this kind loses its low end as resonance rises, because the resonance works by feeding the signal back around the circuit and that feedback subtracts the bass: with the cutoff at 1 kHz and resonance at three-quarters, our ladder measures 13 dB down at 50 Hz. With Bass on, the signal going into the ladder is lifted by exactly what the feedback takes away, so the low end stays where it was however far you turn the resonance up: the same measurement then reads 0.1 dB down instead of 13. The resonant peak still behaves normally, so you keep the sharp, singing resonance without the sound going thin underneath it. It is off by default, so every preset sounds as it did and it is the same idea as the RES COMP switch on the Groove Synthesis 3rd Wave. Expect the filter to sound fuller and a little louder with it on, because the bass that used to disappear is now there. If you switch Bass on and then change the filter to a type that is not a low-pass, the button disappears and Bass does nothing until the filter is a low-pass again.

LP Poles / HP Poles - shown only when MODEL is LADDER and TYPE is MULTI. Together they set the slope in the way made famous by classic American polysynths: the number of low-pass poles and the number of high-pass poles are chosen independently, from none to four each, giving every combination from a gentle single pole to a 24 dB band-pass. LP 4 with HP 0 is exactly the LP tap; HP 4 with LP 0 is exactly the HP tap.

LPG Decay - shown only under the LPG model. How long the vactrol¹ takes to close, in milliseconds. Nothing holds the gate open; it always closes and this is the only control governing how long that takes.

Shift / Mix - shown only under the SHIFT model. Shift sets how far every partial moves, in hertz, up or down; Mix blends the shifted signal against the dry. At a shift of zero the two are the same signal, so Mix has no audible effect until Shift is set away from zero.

Vowel - shown only under the VOWEL model, where it takes the place of the Cutoff knob. It sweeps the filter through seven vowels in the order I, E, AE, OE, A, O, U, which is the order a mouth moves through them from front to back and the knob is labelled with them. The seven sit evenly across the knob's travel, the same at every pitch and the value readout shows the vowel you are on rather than a frequency. Between two vowels the filter is a blend of both.

Voice - shown only under the VOWEL model. The voice part: BASS, TENOR, CTENOR (countertenor), ALTO or SOPRANO. The same vowel sounds different from a bass and from a soprano and not only because of pitch; the strengths of the individual formants change too and that is what this control changes. Between two parts the filter is a blend. The knob steps from part to part but modulation glides smoothly between them, so an LFO or envelope on Voice morphs one voice part into the next.

Formant - shown only under the VOWEL model. Shifts every formant together, up to two octaves in either direction, which is the effect of a longer or shorter vocal tract. It is independent of the note you play, so a vowel keeps its identity while the whole mouth gets bigger or smaller.

1x / 2x - shown only under the VOWEL model, a switch sitting beneath Resonance. It sets how hard the vowel is imposed. At 1x the signal passes through the formant filters once. At 2x it passes through them twice, which doubles every boost and every cut, so the vowel speaks strongly and the sound in between the formants falls further away. 2x is a lot to have on all the time, which is why it is a switch; 1x is the default.

Cutoff - the corner frequency. Under the VOWEL model this knob becomes the Vowel control described above.

Resonance - emphasis at the cutoff, up to self-oscillation (LADDER) or a long resonant ring (ZDF - see below for why these differ). Under COMB it is the feedback of the delay line and under LPG it is a true resonance at the gate's own cutoff. Under VOWEL it sets how narrow each formant is: low values give broad, soft formants, high values give sharp, whistling ones. It has no meaning under SHIFT and greys out there, as do Cutoff and Env Amount - a frequency shifter has no cutoff, no resonant element and no saturation stage.

Env Amount - how much the filter envelope moves this filter's cutoff, bipolar so it can open or close. Under the VOWEL model it moves the vowel position instead, so the filter envelope can sweep the mouth from one vowel to another over the course of a note; at 100 percent the envelope covers the whole set of vowels. The label reads ENV AMOUNT on the panel.

¹ A vactrol is a small electronic part made of a light and a light-sensitive resistor sealed together in a light-proof package. The light is an LED or a tiny bulb, and the resistor is a photocell. When a control voltage lights the lamp, the photocell's resistance falls and more signal passes through. When the lamp goes dark, the resistance rises again and the signal is shut off. The name was originally a trade name of Vactec, the company that made them.

They matter in synthesisers because a photocell is slow and uneven. It responds to the light quickly but recovers slowly, taking longer the darker it gets. A low-pass gate (LPG) uses a vactrol to close a filter and a volume control together. That slow, uneven recovery is what gives the sound its natural, plucked character, like a struck drum or a marimba. Don Buchla made this famous in his 292 module, and it is the "Buchla bongo" sound.

That is why the manual says "nothing holds the gate open; it always closes." The LPG model imitates the photocell letting go after it has been lit, and LPG Decay sets how long that takes.

KB Tracking - keyboard tracking: how much the cutoff follows the played pitch. OFF, 1:2 (half-tracking) or 1:1 (full tracking), so the tone can stay consistent up and down the keyboard. Under the SHIFT model it scales the shift frequency instead, which is the difference between a shift that holds the same musical interval at every pitch and one that tears the bass apart while barely touching the top. A 200 Hz shift with tracking off is more than two octaves at the bottom of the keyboard and a few semitones at the top; at 1:1 it is the same interval throughout. Under the VOWEL model it greys out, deliberately: a real voice does not move its formants with pitch and holding them still is what tells the ear that a bass and a soprano are different sized people rather than one voice transposed.

In the instrument's initial state the two filters are identical: both LADDER, both taking the LP tap, both fully open, both keyboard-tracking at 1:1 - a deliberately neutral starting point from which any patch begins by closing things down rather than hunting for why the top end is missing. Each oscillator chooses which filter it feeds: in the initial state OSC 1 and OSC 3 feed Filter 1 while OSC 2 feeds Filter 2, so the two default oscillators arrive on independent filters from the first note.

Two further controls govern how the filters work together. Link Filters couples the two cutoff controls so that adjusting one moves the other, letting you sweep both as a pair while keeping their offset. Filter Order sets the routing between them: Series runs the signal through Filter 1 then Filter 2 for steeper, combined shaping; Parallel - the initial setting - splits the signal through both and sums the results for blended responses such as a low pass plus a high pass; and Stereo places one filter on each channel for independent left/right tone.

11.3.5.1 THE FILTER MODELS

The Model selector chooses between six circuits - not tone presets or EQ tilts but genuinely different pieces of mathematics with audibly different personalities. Two are classic filters and four are not filters in the ordinary sense at all. All are fed identically: the cutoff you set, the envelope, keyboard tracking, the modulation matrix and the smooth per-block coefficient ramps all compute once, upstream of the model and deliver the same values to whichever circuit is selected. Switch models and only the circuit changes; every modulation behaves the same. The one exception is VOWEL, which reads the cutoff knob as a vowel position rather than a frequency, as described below.

LADDER is a four-pole transistor ladder and it is modelled as one rather than as an idealised filter. In the original circuit every stage is a pair of transistors and a transistor pair does not pass a signal in a straight line: it compresses the peaks, softly and more so the harder it is driven. NeuralDream reproduces that in every stage and in the resonance path too, which is why this filter has a body and a bloom that a textbook filter does not. Play it quietly and it is warm; push the mixer into it and it thickens and growls before it ever sounds distorted, exactly as the hardware does. The low pass is a true 24 dB per octave; the band pass and high pass are not separate filters but are drawn from the same ladder by mixing its four stages, so all three share the family character. Resonance at full settles into a clean, sustained whistle that the transistors hold at a steady level and the passband loses bass as resonance rises, which is a characteristic of the original and is left as it is. This is the filter that gives the synthesiser its voice and the reason LADDER is the default. Appendix A describes how it is made.

ZDF is a state-variable filter in the trapezoidal-integration form popularised by Andrew Simper (Cytomic) - the same Zavalishin lineage, an entirely different topology. Its outputs are 2-pole, so low pass, high pass and band pass are 12 dB/octave: half the slope, twice the gentleness. Resonance produces the SVF's characteristically clean, symmetric peak and at maximum it behaves differently from the ladder in a way worth understanding: where the LADDER truly self-oscillates - a sustained sine that sings indefinitely - the ZDF at full resonance is an extremely high-Q ring that decays over tens of milliseconds, like a struck tuning fork rather than a bowed one. Pluck it with an envelope or a note-on transient and it chimes; it will not drone. This is faithful to the algorithm's nature and it is the same state-variable core used by the sampler engine's filter, so a patch moved between the two engines with matching settings now has a matching filter available.

COMB is a delay line with feedback - Cutoff sets the delay length and therefore the pitch of the comb; Resonance sets the feedback and therefore how long it rings. It has no taps, so the Type control greys out. The Comb type stands apart from the model choice entirely. In the synthesiser it is always the Karplus-Strong feedback comb (Karplus and Strong, 1983): a delay line whose length is pinned to the cutoff knob so the comb's fundamental plays the cutoff frequency, with a 2-sample moving-average damper in the feedback loop - the classic plucked-string recipe in which high frequencies decay faster than lows, so held combs mellow the way a real string does. Selecting ZDF does not substitute a different comb; the sampler engine's Hermite-interpolated comb remains the sampler's own.

LPG is a low-pass gate, the West Coast idea in which a resonant two-pole low-pass whose cutoff and level both fall together, driven by a vactrol that is struck at note-on and closes on its own. That falling-together is what separates a gate from a filter - a note does not merely get quieter as it decays, it gets darker at the same time, which is what makes a struck object sound struck. LPG Decay is the only control over how long that takes; nothing holds the gate open. Cutoff sets how bright the strike is and Resonance is a real resonance at the gate's own cutoff.

SHIFT is a Bode-style single-sideband frequency shifter and it is not a filter at all. Where a filter attenuates, a shifter moves every partial by the same number of HERTZ rather than the same ratio - so a harmonic series does not stay harmonic. Small amounts give slow beating and phasing; large amounts turn tones into bells and destroy any sense of a fundamental. Only Shift, Mix and KB Tracking apply; Cutoff, Resonance, Drive and Env Amount all grey out.

A note on the shifter's image rejection, since it is audible at the bottom of the range: the unwanted sideband is suppressed by about 80 dB above 4 kHz but only around 8 dB near 440 Hz. This is inherent - the mathematics a shifter depends on is discontinuous at zero hertz and no finite filter approximates it there. A Bode analogue shifter has exactly the same weakness. A trap below 30 hertz removes the sub-audio energy that would otherwise accumulate and eat headroom.

VOWEL is a formant filter and it is not a filter in the ordinary sense either. Where a filter has one cutoff, a vowel is defined by five resonant peaks at once, the formants, whose frequencies, strengths and widths are what make an "ah" an "ah" and an "ee" an "ee". The model holds a table of those five formants for each of seven vowels and each of five voice parts, measured from real singers and builds five resonators in parallel from it every time a control moves. Turning the Vowel knob moves smoothly through the table, so between "A" and "O" you hear a mouth in transition rather than a switch. The Voice control chooses which singer's table is used and Formant scales all five peaks together for a bigger or smaller mouth. Resonance sets how narrow the peaks are. A gentle tilt in front of the resonators stands in for the natural fall of the voice's source, so the model responds best to a source with plenty of harmonics: a saw, a rich wavetable or the Partial engine. The 1x / 2x switch beneath Resonance sets how hard the vowel is imposed; 2x doubles every boost and cut. The filter's loudness is held steady across the whole vowel sweep and across the voice parts, so moving between vowels changes the colour and not the level.

Choosing between them is a matter of role. Reach for LADDER when the filter is the voice of the patch: basses that need weight, leads that need the resonant whistle, anything where a steep sweep should feel like an event. Reach for ZDF when the filter is a lens rather than a voice: taming a bright wavetable without amputating it, gentle high-pass cleanup that leaves the body intact, resonant plucks that chime and step aside or surgical band-pass work where the SVF's symmetry keeps the result transparent. With two filters, six models and three routings, the combinations are their own instrument: a ZDF high pass clearing room in series before a LADDER low pass does the singing is a mixture no single filter architecture could give you and a VOWEL on Filter 2 in parallel with a LADDER on Filter 1 puts a mouth on a bass without losing the bass.

Four ADSR envelopes shape every voice. The VCA Envelope controls the voice's amplitude - its attack, decay, sustain and release define the overall shape of the note from key-down to the end of its release tail. The Filter Envelope is a second, independent ADSR whose primary job is to move filter cutoff over time, by the amount each filter's Env Amount control allows. Two further envelopes, Env 3 and Env 4, have no fixed destination - they exist purely as extra per-voice modulation sources for the matrix (see 11.4.2).

The initial state gives both main envelopes musical proportions rather than test-bench minimums: decay of 2 seconds and release of 2.5 seconds on the VCA and Filter envelopes alike, so a fresh INIT plays a note with natural length and tail from the first key. Double-clicking any envelope knob returns it to these same values.

All four envelopes share the same extra options. Velocity sets how strongly note velocity scales the envelope, so harder playing can mean louder notes or a brighter filter sweep. Loop cycles the envelope's attack and decay stages continuously while a note is held, turning the envelope into a repeating modulation source. Sync locks that looping to the host tempo instead of free-running: with Sync on, the loop restarts on a musical division rather than at a rate in seconds, so the pulsing stays in time with the transport - the same Sync behaviour the LFOs offer.

11.3.7 THE LFOs

Four low-frequency oscillators provide cyclic modulation. They are per-voice, so each note runs its own set and they offer the following controls:

Wave - Sine, triangle, saw, ramp, square, sample-and-hold or random.

Rate - Free-running speed, in hertz.

Sync / Sync Rate - When Sync is on, the LFO locks to the host tempo and Rate is replaced by a musical division (such as 1/4 or 1/8).

Fade - A fade-in time, so the LFO eases in after the note starts rather than appearing at full depth immediately.

Retrig - When on, the LFO restarts its cycle at each new note; when off, it free-runs continuously so notes catch it at different points in its phase.

The LFOs do not connect to anything on their own - they are sources for the modulation matrix, described in Section 11.4.

11.3.8 GLOBAL PERFORMANCE CONTROLS

Glide Portamento - the time taken to slide pitch from one note to the next, from zero (off) up to two seconds.

Bend - The pitch-bend range, in semitones, for ordinary channel pitch bend (the default is two semitones).

MPE Bend Range - The per-note bend range for MPE member channels, set in the configuration area (see 3.8). The default is 36 semitones.

Analog - Voice-to-voice drift and variance, as described in 11.3.1.

11.3.9 MPE AND EXPRESSIVE CONTROL

The synthesiser is fully MPE-aware (MIDI Polyphonic Expression), which lets controllers such as the Expressive E Osmose, ROLI Seaboard and LinnStrument shape each note independently after it has been played. Three dimensions are handled per note:

Per-note pitch bend - each note can bend on its own member channel, independently of the others, over the range set by MPE Bend Range. With the default of 36 semitones a single note can be slid three octaves, which suits the continuous-pitch playing style of the Osmose and similar instruments.

Pressure - continued pressure on a held note (channel pressure on a member channel or polyphonic aftertouch). It is available as a modulation source.

Slide - the timbre or vertical dimension (MIDI CC 74). It is also a modulation source.

The engine adapts automatically. On a true MPE controller, Pressure and Slide are read per note, so each voice responds to its own movement. On an ordinary single-channel keyboard the same

sources still work, applied globally, so the same routings remain useful whatever you are playing. The MPE member bend range is set in the synthesiser's configuration area, opened from the cog control; the value is saved with the preset.

11.4 MODULATION

The synthesiser has its own per-voice modulation matrix. This is separate from and complementary to, the global modulation

strip described in the main manual and the distinction matters.

11.4.1 PER-VOICE MODULATION VS THE GLOBAL STRIP

The global modulation strip - the four LFOs, four envelope followers and four macros - produces one value that is shared by the whole plugin at any instant. Routed to a synthesiser control, it moves that control identically for every voice at once. That is exactly what you want for a unified sweep across a chord.

The synthesiser's own matrix is different: it is per voice. Each note runs its own LFOs and envelopes, so a route inside the synthesiser matrix moves each note independently. A filter envelope routed to cutoff opens a fresh sweep for every note as it is played; an LFO routed to Position gives each note its own drifting timbre. This per-note independence is what makes the result feel like a real instrument rather than a single sound being processed.

Both are available at once. Synthesiser controls are ordinary modulation targets, so the global strip can reach them as well - giving you a shared, chord-wide movement and an independent per-note movement on the same control if you wish.

11.4.2 THE PER-VOICE SOURCES

The synthesiser matrix offers a full set of per-voice sources. Every one is evaluated independently for each note, so a chord's voices modulate independently of one another.

LFO 1-4 - the four per-voice LFOs (see 11.3.7).

Filter Envelope and VCA Envelope - the two primary ADSRs (see 11.3.6), usable as modulation sources in their own right as well as in their primary roles.

Env 3 and Env 4 - the two spare ADSRs, which exist solely as modulation sources.

Pressure and Slide - the MPE expression dimensions (see 11.3.9), per note on an MPE controller and global otherwise.

Velocity - the note's played velocity, latched at note-on.

Random - a fresh random value per note, latched at note-on (a per-voice sample-and-hold).

Key Track - the note's position on the keyboard, bipolar, centred on middle C.

Lift - the note's release velocity, latched at note-off.

SM 1-4 - the four sequencer modulation lanes (see the Step Sequencer section). Each outputs its current step value; they follow the sequencer clock and the RETRIGGER switch and they run whenever the transport is playing even with the sequencer's note lanes disabled.

11.4.3 CREATING A ROUTING

Routings are created in the same way as anywhere else in NeuralDream: right-click (or control-click) the knob or slider you want to move. This opens the MODULATION overlay for that control. Add a source, set its amount (depth) and choose whether it is unipolar or bipolar - unipolar moves the control in one direction from its current setting, bipolar moves it equally above and below. Each route can also be given a VIA source, a second source that scales the route's depth so the modulation itself swells, gates or fades under another control (see Section 11.4.3.2). You can add several sources to one control; their contributions sum. All routings, including their VIA assignments, are saved with the preset.

Nearly every control on the synthesiser can be a per-voice target and as of v2.13 that includes every control on the GRAINS panel, the STRETCH, SHIFT, WIDTH and AIR controls on the PROCESS tab and DETUNE. As of v2.17.31 it also includes TILT on the PROCESS tab, the filter Voice, the DELAY on each envelope, Bend and UNISON. UNISON is set at the moment you press a key, so each new note takes the knob plus the modulation at that moment and the count cannot change while the note is held. A source that has a value at the key press, such as Velocity, Key Track or Random, gives each note its own unison count. A few controls are not offered as targets, either because they are global to the instrument rather than belonging to a voice or because they are settings rather than things to sweep: the master output level, the SMOOTH control on each LFO, the Starling settings, ODD/EVEN on the PROCESS tab, LP Poles and HP Poles on the filters and NOTE OFFSET on a Modal oscillator. If you cannot find a control in the target list, that is why. Note also that AIR at zero is a switch rather than a level: the engine does not generate its noise at all when the knob is off, so modulation moves the amount once the knob is up but cannot switch AIR on from zero.

11.4.3.1 Q AND P - STEPPING A MODULATION

Every route carries two buttons that turn a smooth modulation into a stepped one and only one of them can be active at a time. Q cuts the route's travel into equal steps. Press it to cycle through 2, 4, 8, 16 and 32 and press it once more to switch it off. With Q set to 4, a mod wheel sweep moves the target in four equal jumps instead of a slide and on a tempo-locked grain loop those steps land on beat boundaries, which is where this began.

P does the same job musically. Instead of equal steps it snaps the route to the notes of a scale, so the target lands only on those notes as the source moves. Press P and a list of scales appears: Chromatic, Major, Minor, Dorian, Phrygian, Lydian, Mixolydian, Pent Maj, Pent Min, Blues and Whole tone, with Off at the top to switch it back off. The button then shows the scale you chose. Beside it is a button reading 1 OCT, which says how far full modulation reaches; press it to cycle through 2, 3 and 4 octaves. It is greyed while P is off, because it has nothing to act on.

The obvious use is pitch. Route the mod wheel to an oscillator's FREQUENCY, set P to Minor and 1 OCT and pushing the wheel from bottom to top walks the oscillator up the minor scale, landing on the root, the second, the minor third, the fourth, the fifth, the minor sixth, the minor seventh and the octave and nothing in between. At 2 OCT the same sweep covers two octaves of the same scale. On a pitch control an octave always means twelve real semitones whatever the control's own range. FREQUENCY reaches 36 semitones either side of its setting. At 3 OCT the sweep fills it exactly and at 4 OCT it stops at those same 36 semitones. A bipolar route steps the same way below the control's setting as above it. FINE reaches only one semitone either side of its setting. P on FINE can therefore land only on the notes of the scale that lie within that semitone. On a bipolar route Chromatic gives one semitone down, the setting and one semitone up. Major and Lydian give one semitone down and the setting. Phrygian gives the setting and one semitone up. Every other scale leaves FINE where it is.

P is not restricted to pitch. It works on any destination, so a filter cutoff stepped by a pentatonic scale moves in intervals that sit musically against what you are playing and there is nothing to stop you putting it on pan if the result appeals. P lands on the notes of its scale only when the route's depth is at 100%. A lower depth shrinks every step by the same amount and the steps then no longer fall on the scale's notes. At 50% depth a minor scale over one octave gives steps of 0, 1, 1.5, 2.5, 3.5, 4, 5 and 6 semitones, several of which sit between two notes. Keep the depth at 100% and choose the range with the octave button when you want every step in tune. Q and P are saved with the preset and travel with a route when you copy and paste it.

11.4.3.2 VIA - MODULATING THE DEPTH OF A MODULATION

A modulation route can now carry a VIA modifier: a second modulation source that scales the *depth* of the route. Where a normal route applies a fixed amount of a source to a target, a route with a VIA source applies an amount that is itself moved by the VIA source - so the modulation can be brought in and out, gated or swelled by something else.

The classic use is expression: route an LFO to a filter cutoff, then set the VIA of that route to Pressure and the LFO wobble only appears - and deepens - as you lean into the note. Any modulation source can act as the VIA modifier and VIA is set in the same MODULATION overlay used to create the route, per route, so different sources on the same control can each have their own VIA. Two entries at the top of the VIA menu, SMOOTH and SLOW, are not sources at all: they change the shape of the route's own signal in time and are described in the next section. As with all routings, VIA assignments are saved with the preset.

11.4.3.3 SMOOTH + SLOW - SHAPING MODULATION IN TIME

The first two entries in a route's VIA menu, SMOOTH and SLOW, do not scale the route's depth by a second source. Instead they pass the route's own signal through a slew before it reaches its target, so a source that moves suddenly arrives as a curve. When either is chosen, the route's VIA amount slider stops being a depth and becomes a time: hover over the slider and its tooltip reads that time in seconds.

SMOOTH treats the rise and the fall alike. Whatever the source does, the target follows it with the same lag in both directions, so an envelope follower's spike reaches a filter as a swell and lets go as a fade and a square LFO's edges arrive as ramps. At short times it takes the click out of a hard edge; at long times it behaves as an integrator and the target follows the average of the source rather than its shape. It is meant to take out the human wobble in a held gesture, not to slow the gesture down, so its range is short: the slider runs from 0 to half a second on a square law, a quarter of the slider being 0.03 s, half 0.125 s and full 0.5 s. Choosing SMOOTH from the menu starts it at a quarter. At zero it is switched off and the route passes through untouched.

SLOW is a capacitor behind a diode: it charges quickly and discharges slowly. The rise follows the source almost immediately, within 10 ms, so a source that arrives fast still arrives fast. The fall is where SLOW acts and the slider sets how long it takes: 4 seconds times the slider squared, so a quarter of the slider is 0.25 s, half is 1 s and full is 4 s. Choosing SLOW from the menu starts it at half. It was built for polyphonic aftertouch. Pressure on a key can be pushed in gradually but the moment the finger lifts it falls from full to nothing at once and anything it is driving snaps with it. Route pressure through SLOW and the pressure still lands as fast as you apply it but when you let go the target eases down over the time you set, as if the note were still ringing under your finger.

Both keep their own memory per route and in the synthesiser per note as well, because a source such as pressure or an envelope differs from note to note. Both are saved with the preset as part of the route.

11.4.3.4 TEMPO AND TIME SIGNATURE

The host tempo and time signature are shown live at the bottom-left of the window, in the strip directly above the input and output meters - for example 120.0 BPM 4/4. They read straight from the host transport and update as the tempo or meter changes and they are the same values used everywhere tempo matters in the synthesiser, including the grain-size bars/beats readout and tempo-synced grain spawning.

11.4.4 COPYING AND PASTING ROUTINGS

The MODULATION overlay has Copy and Paste buttons in its header. Copy captures the complete set of routings on the current control - every source, its amount and its unipolar or bipolar mode. Right-click another control and press Paste to apply that same set to it, which makes it quick to give several controls the same modulation. The buttons briefly highlight when clicked so you can see the action has registered.

Copy and Paste also work between the synthesiser and the ordinary effect engines. There is one sensible limitation: the per-voice synthesiser sources (the synthesiser's own LFOs and envelopes, Pressure and Slide) exist only inside the synthesiser. If you copy a control that uses them and paste onto an effect-engine control, those per-voice sources are dropped and only the global sources (the global LFOs, envelope followers, macros and MIDI sources) are pasted across. Pasting between two synthesiser controls keeps everything.

11.4.5 WORKED EXAMPLE: PRESSURE VIBRATO

Nothing teaches the modulation system faster than building the patch every expressive player asks for first: vibrato that lives under your finger. The goal - press into a note and vibrato blooms on that note; ease off and it stills; and on an MPE controller, each note answers only to its own finger while its neighbours hold steady. One route builds it and the shape of that route is a lesson in how the whole matrix thinks.

Open the modulation panel on the oscillator's FINE control and add Synth LFO 1 as a source, with a small depth - around +10% is a musical starting point. This alone gives constant vibrato: the LFO wobbles the fine pitch of every note, always on. The expressive half is the VIA: set the route's VIA to Pressure and push the VIA amount to full. VIA is a second source that scales the first route's depth - so the LFO still supplies the wobble but Pressure now supplies how much of it each note receives, from nothing at rest to the full $\pm 10\%$ under firm pressure. Source makes the movement; VIA decides who gets it and how much. That division of labour is the heart of the matrix.

Why this is truly polyphonic on an MPE controller comes down to which Pressure you use. The synthesiser's per-voice source list includes Pressure and it is per-note by design: an MPE instrument sends each sounding note on its own MIDI channel and every voice reads the pressure of its own channel only. Press into one note of a held chord and only that voice's route deepens. (The instrument also has a global Aftertouch among its rack-level modulators - the right tool for whole-patch gestures and the wrong one here: a route or VIA built on the global source moves every note at once. If your "per-note" vibrato ever behaves globally, this substitution - or routing Pressure at the LFO's own depth control, which is a single shared parameter - is almost always the cause.)

11.5 MAXIMISING THE WAVE DISPLAY

Each synthesiser wave display has a small maximise button (beside the cog, in the display's top-right corner). Pressing it pushes the oscillator's wave display up into the large visualiser area at the top of the plugin, giving a big, detailed view of the wavetable, grain cloud or partial spectrum. Press it again to return the display to its normal size. If the visualiser is folded away (see Section 8.2), pressing it opens the visualiser and switches XFM on; folding the visualiser away again returns the display to its normal size. While maximised it continues to animate live.

In the maximised view of a sample-based oscillator the mouse wheel becomes a zoom and navigation tool. Vertical scrolling zooms about the cursor, from the full sample down to better than a thousand-fold magnification; horizontal (side) scrolling or shift-scrolling, pans the view left and right through the sample. Wheel input over the display belongs to the display: it zooms and pans the waveform only and never scrolls the plugin window or the layer stack beneath - the same containment applies over every engine visualiser.

The wavetable displays work the same way, at normal size as well as maximised and so does the PARTICLES visualiser. Scrolling sideways slides the view along the single cycle; scrolling up and down zooms in and out about the pointer, so you can look closely at one part of the waveform; shift with the wheel slides the view as well. Every frame of the table is still drawn while you are zoomed in: only the part of the cycle in view changes. Double-click puts the zoom back to the whole cycle and in the three-dimensional view it returns the viewing angle to straight on at the same time.

Each oscillator keeps its own view. Zooming or turning the display for oscillator 1 leaves oscillators 2 and 3 as you left them. Each one is shown its own way again when you select it. Each PARTICLES layer keeps its own view in the same way. A view stays as you set it while the plug-in is open, even if you close and reopen its window. Views are not saved in presets or in your project: INIT, reopening a project and a new instance of NeuralDream all start every display at the whole cycle, straight on.

As you zoom, the rendering moves through three regimes, each drawing the truth available at that magnification. At full view you see the familiar envelope outline of the whole sample. Once meaningfully zoomed, the display switches to true per-pixel minimum/maximum columns computed from the samples actually visible - the treatment a DAW audio editor uses - so transients, asymmetry and waveform shape appear instead of a stretched summary; this view is cached against the current window, so a static view costs nothing to animate. Zoom further still and the display hands over to a sample-accurate rendering: the actual waveform polyline and, at the highest magnifications, the individual sample points themselves, drawn over a brightened zero line. This is the view for precision work - zero crossings are directly visible, which pairs naturally with the granular engine's own zero-crossing spawn alignment: you can see exactly the crossings the engine will favour and place OFFSET, INSET and POS against them by eye.

Two further view toggles sit in the tab row beneath the wave display and apply to wavetable oscillators only (they are hidden in Granular and Partials modes, where they have no meaning):

2D/3D - switches the wavetable display between a flat 2D view of the current cycle and a 3D "waterfall" view of the whole table, so you can see the frames the Position control sweeps through. Simple single-cycle shapes (Sine, Saw, Triangle and the rest) are shown in 3D as a notional wavetable whose slices are the shape at every width the Position control gives it, from the complete shape to silence. You see the widths POS sweeps through just as you see the frames of a table.

OSC/SUM - switches the display between showing the focused oscillator alone and the summed output of all the oscillators, so you can view one voice's shape or the combined result.

How the three-dimensional display is drawn

In the three-dimensional view the wavetable is drawn as a landscape running away from you. Every single frame of the table is drawn, however many it holds, each as one thin line: a table of a thousand frames shows a thousand lines, not a sample of them. The lines nearest you are the brightest and they fade gently into the distance, which is what gives the display its depth. Behind them the backdrop runs from black at the top to a deep blue at the bottom, so that even the faintest lines at the back stay visible.

The frame the oscillator is actually playing is drawn in white on top of the landscape, so you can follow it as POS moves through the table. Drag inside the display to turn the landscape, drag up and down to tilt it and double-click to put it back to straight on.

Use your mouse scroll-wheel or apple mouse touch scroll up and down to increase and decrease magnification and left and right to scroll within the Wave Display - this also works when you maximise the Wave Display to the Main Visualiser and maximise ([]) that to the entire plugin's window.

The PARTICLES visualiser draws its table with exactly the same code, so the two displays always look and behave alike. Only its colour differs: its lines are drawn in the particles orange, with the frame WT POS is reading in white.

11.6 GLOBAL SETTINGS

NeuralDream now keeps a small set of global settings that persist across every instance of the plugin and every project, stored in a settings file on disk separate from the per-preset state your host saves:

- **macOS:** `~/Library/Application Support/DreamDeep/NeuralDream.settings`

- **Windows:** `%APPDATA%\DreamDeep\NeuralDream.settings`

These act as the default that individual presets fall back to. The instrument and the effect each keep their own set inside that one file, so the two never overwrite each other when both are loaded on the same track: you can give the instrument one backdrop and the effect another, hold the modulation strip open in one and closed in the other and cap their voices differently. The single setting they share is the author name used when you save a preset, because that describes you rather than a plugin. The primary setting is a voice cap (a ceiling on synthesiser polyphony), offered as Mono, 2, 4, 6, 8, 12, 16, 24 or 32 voices, with 16 as the factory default. A user on a lower-powered machine can set, say, a six-voice cap once and every preset and every new instance will honour it without having to edit each sound - while a preset can still specify its own value where it needs to, from the same list.

Global settings are reached from two places. The cog in the preset strip opens Display Settings, which carries the settings that apply to the whole plugin regardless of what is loaded - the background and the modulation strip's visibility. Because these need no synthesiser, that dialog is available in every version of the plugin and in every preset and what you set there applies to the plugin you set it in, leaving the other one as it was. The cog in the synthesiser's wave display opens the synthesiser's own configuration - MPE range, velocity curve and the voice cap - where each setting is offered as a global default and, where appropriate, a per-preset override.

11.7 - STEP SEQUENCERS

The Instrument has two step sequencers. Each plays looping patterns for gate, pitch and velocity whenever you hold a key, each has its own rate, swing, direction and length and each can be sent to different oscillators, so one sequencer can drive a bass line on oscillator 1 while the other runs a faster figure on oscillators 2 and 3. Both are always synchronised to the DAW's tempo and song position and both drive the synth by generating notes exactly as if they had arrived over MIDI - every sequenced note passes through the same note-on/note-off path as a note you play yourself, so envelopes, glide, voice legato and modulation all behave identically. The two sequencers are edited one at a time in the same panel; Section 11.7.12 explains how to switch between them and how to decide which oscillators each one plays.

11.7.1 OPENING AND ENABLING

SEQUENCER button - sits in the preset strip, to the right of the cog button. Click it to switch the main visualiser into the sequencer editor; click again to return to the waveform view. If the visualiser is folded away (see Section 8.2), it opens and XFM switches on; folding the visualiser away again closes the editor. The button lights while the editor is shown. The sequencer view and the maximised waveform view are mutually exclusive - opening one closes the other.

STARLING button - sits immediately to the right of SEQUENCER and opens the Starling view in the same space (see Section 11.8). If the visualiser is folded away (see Section 8.2), it opens and XFM switches on; folding the visualiser away again closes Starling. SEQUENCER and STARLING are two views of one area: click either to switch straight to it or click the one already showing to close it.

PRESETS - sits at the end of the row of buttons beneath the synthesiser's wave display, to the right of the keyboard icon. It opens the same preset list as before.

SEQUENCER enable - the yellow dot at the top of the right-hand column turns the sequencer on and off. This is separate from displaying it: you can edit a pattern with the sequencer disabled and you can close the editor while the sequencer keeps playing.

SEQ 1 / SEQ 2 chip - sits in the top row of the panel, just to the left of the three oscillator squares. It tells you which of the two sequencers you are looking at. Click it to switch to the other one: the tabs, the steps, the right-hand column and the playhead all change to show that sequencer and the chip turns copper and reads SEQ 2 while the second sequencer is on screen. Each sequencer has its own enable dot, so either, both or neither can be running regardless of which one is on screen.

The sequencer plays whether or not the DAW transport is running. Stopped, it free-runs at the host tempo and always plays phrases from step 1, per key; playing, it locks to the song position.

11.7.2 PLAYING

Hold a key and the pattern plays, transposed to that key. Hold a chord and every held note runs the pattern in parallel. Release the key and its sequenced notes stop, exactly as if you had stopped playing them.

Launch quantise - sequenced notes only ever trigger on the grid set by the RATE control, regardless of when you press the key. Play slightly early and the phrase waits for the next RATE boundary; press keys at slightly different moments within the same step and their phrases share the same boundary and trigger in exact lockstep, each running from its own step 1. In timeline mode (RETRIGGER off) everything is grid-locked to the arrangement by definition.



11.7.3 LANES

The pattern is made of colour-coded lanes, selected by the tabs across the top. The 16-step display always shows the currently selected lane, repainted in that lane's colour. Sequencer 1 has all seven lanes listed below. Sequencer 2 has the first three only - GATE, PITCH and VEL - because the four SM modulation lanes belong to sequencer 1 alone, so that there is only ever one SM1, one SM2 and so on to route.

Tab	Colour	Controls
GATE	pink	Note length per step, legato, on/off
PITCH	teal	Semitone offset per step, $\pm$ octave
VEL	gold	Note velocity per step
SM1	violet	Sequencer modulation lane (see 11.7.6)
SM2	orange	Sequencer modulation lane (see 11.7.6)
SM3	green	Sequencer modulation lane (see 11.7.6)
SM4	blue	Sequencer modulation lane (see 11.7.6)

Step controls (every lane):

Playhead - the currently sounding step is highlighted (lit column with a bright marker beneath it). Each lane shows its own position, so with unlinked lengths polyrhythmic lanes visibly cycle at their own rates. With nothing held, the playhead follows the running clock (the SM lanes are always live).

Step number (top) - click to toggle the step on/off. A disabled step is dimmed and muted: no note plays and the position is skipped silently.

Slider - drag or click to set the step's value. Movement is smooth and continuous unless the lane's Q setting quantises it.

Value readout - beneath each slider, in the lane's units: gate %, pitch in semitones, velocity 0-127.

CHANCE row - the bottom number row on every lane. 1 (dim) means the step always plays; drag it up to N (lane colour) for a 1-in-N chance, up to 100. Each step rolls a fresh, independent random dice every time it comes round, so a step set to 3 plays roughly one pass in three with no fixed pattern - it never falls into a repeating rhythm, even under a DAW loop. A muted step ignores chance entirely. On the gate lane a failed roll skips the whole step (ratchets included); on pitch it plays the held note unshifted; on velocity it falls back to your played velocity; on SM lanes it outputs 0.

Aux row (between the value readout and CHANCE) - on the GATE lane this is the ratchet: drag vertically to make the step fire 2-6 times, evenly dividing the step (a dot means a single hit). On the PITCH lane the same row holds OCT (± 3 octaves). Ratchet is meaningful only on the gate lane - sub-hits re-read the same pitch/velocity step - so the two controls share the row without ever overlapping.

GATE lane - the slider sets note duration per step, low for short staccato notes, high for long ones. Legato: drag the slider above the top of its range - the readout shows LGT with a bracket, the step plays full length and the next step's note overlaps it instead of retriggering, so the synth's own Voice legato and glide carry the envelopes across the join. For authentic legato/acid lines, enable Legato in the Voice section and add some glide.

PITCH lane - values are bipolar offsets from the incoming note: centre (0) plays the held key unchanged, +2 plays two semitones up, negative values go down, range ± 12 semitones, shown signed. With Q OFF the lane is truly continuous (the readout shows one decimal, e.g. +2.4) and fractional pitches are delivered as per-note MPE pitch bend - each fractional note takes its own member channel with bend sent before the note-on, exactly as an MPE controller would play it; integer pitches use no bend. If you also play a live MPE controller, the sequencer's member channels can overlap with the controller's - an accepted trade-off of the pure-MIDI design. The OCT row adds up to ± 3 octaves per step. Set Q to 12 for exact semitone steps (see 11.7.4).

VEL lane - the slider sets each step's note velocity (0-127). For this to be audible the patch must respond to velocity - raise the VEL amount in the Filter and/or Amp envelopes.

11.7.4 Q - PER-LANE QUANTISE

The Q pill sits beside the tabs, drawn in the current lane's colour: it belongs to the selected lane only and every lane has its own setting.

Q OFF - values are smooth and continuous.

Q 1-16 - snaps the slider to that many divisions; the bar and readout snap visibly while you drag and the underlying fine position is remembered (set Q back to OFF to recover it). On the PITCH lane, Q is divisions per octave: Q 12 = exact semitones, Q 6 = whole tones, Q 4 = minor thirds, Q 1 = octaves.

Click the left half of the pill to decrease, the right half to increase.

11.7.5 LANE LENGTH AND LINK

Length slider (bottom) - drag to set how many of the 16 steps the lane plays (1-16).

L icon / LINK LANES - toggles linked lengths. Linked (slider drawn white): one length for all lanes, dragging on any tab moves them all together. Unlinked (slider drawn in the lane's colour): each lane has its own length. This is where polyrhythms live - GATE at 16 against PITCH at 8 makes the melody cycle twice per rhythm cycle; 16 against 12 or 5 drifts and realigns over longer spans.

11.7.6 SM1-SM4 - SEQUENCER MODULATION SOURCES

The four SM lanes are step patterns that act as modulation sources, exactly like the LFOs: right-click any modulatable knob and they appear in the per-voice section of the source list as SM 1-4. Each lane outputs its step value (0-1); a muted step outputs 0.

They follow the sequencer clock and the RETRIGGER switch: with retrigger on, every voice reads the lane at its own phrase position; with retrigger off, all voices share the DAW-timeline position.

They run whenever the transport is playing, even with the sequencer disabled - so they double as free-running step-LFOs without sequencing any notes.

Each SM lane has its own length and Q, so an SM lane at a different length from the note lanes makes evolving polyrhythmic modulation.

11.7.7 GLOBAL CONTROLS (RIGHT-HAND COLUMN)

Control	Function
SEQUENCER	SEQUENCER - Enable/disable (yellow dot). Each sequencer has its own dot, so either, both or neither can be running
SEQ 1 / SEQ 2	Shows which sequencer is on screen; click to switch to the other. See Section 11.7.12
OSC 1 2 3	Three small squares that decide which oscillators this sequencer plays. Lit means this sequencer drives that oscillator. See Section 11.7.12
M	Lane mute, beside Q. It belongs to the lane whose tab is selected. A muted lane is bypassed entirely and behaves as if it were flat: a muted GATE lane fires every step at half length, a muted PITCH lane plays the held key unshifted, a muted VEL lane uses your played velocity and a muted SM lane outputs zero
LATCH	When on, the pattern keeps playing after you lift the keys, exactly as it played while you held them. Press a new key or chord after everything has been lifted and it replaces what was latched rather than adding to it. Switch LATCH off and any latched notes are released
ARP	Chord mode. OFF plays the pattern once per held key, in parallel. UP, DOWN, UP-DOWN, PLAYED and RANDOM instead treat the held keys as one chord and play the pattern through its notes in that order, one note per step, exactly as an arpeggiator does
OCT	Beside ARP: how many octaves the arpeggio spans, 1 to 4, before it returns to the bottom
RATE	The step size as a division of the DAW tempo (click left/right halves to step through the divisions)
SWING	Timing feel. 0% (up to 50%) = straight; 66% = traditional 2:1 shuffle; higher pushes harder
GATE SC	Gate Scale - scales all gate-lane lengths, 100% = as programmed
DIR	Playback order (for a 4-step lane): Forward 1-2-3-4 · Reverse 4-3-2-1 · Ping-Pong 1 1-2-3-4-3-2 (ends not repeated) · Ping-Pong 2 1-2-3-4-4-3-2-1 (ends repeated) · Random
LINK LANES	Same as the L icon by the length slide - makes all lanes the same length
LOAD / SAVE	Save or load the complete setup of the sequencer currently on screen - every lane, every value and all its globals - as a .seq file in your presets folder, independently of presets. SAVE prompts for a name and asks before overwriting; LOAD shows an alphabetical list of every .seq file found. Loading replaces the entire pattern of the sequencer on screen and leaves the other sequencer untouched, so a file saved from sequencer 1 can be loaded into sequencer 2. A starter pack of melodic factory sequences (major/minor noted in each name) is available - drop the .seq files into your presets folder and they appear in LOAD
RETRIGGER	ON - the pattern restarts from step 1 with every key you play. OFF - the pattern position follows the DAW timeline, so keys played mid-bar join the pattern mid-flight. With the transport stopped, both modes play phrases from step

Control	Function
QUANTISE IN	ON: every key you play waits for the next step at the sequencer's RATE, whether it starts the sequence from silence or joins one that's running. However loosely you play, everything lands on the beat. A key released before its step arrives is not heard. With LATCH also on, a new key does not interrupt the latched sequence: the latched keys keep playing until the new key's step and the new key takes over exactly on that beat, so a change of key is seamless. OFF: every key starts its own sequence the moment you press it, including keys added while others are held. Each one runs on its own start-point, so off-beat playing sounds off-beat - use this if you have perfect timing or want to introduce movement perhaps in some kind of soundscape. Two things are not affected: in ARP mode the arpeggio remains one shared pattern and with RETRIGGER off while the DAW is playing, every key follows the DAW timeline, so a new key sounds the current step straight away rather than starting a sequence of its own. QUANTISE IN is on by default, including in presets saved before it existed and it is saved with each preset and each .seq file.

11.7.8 SAVING AND RECALL

Everything is stored automatically for both sequencers - every step of every lane, all lane lengths and Q settings, all globals including whether each sequencer is enabled, which oscillators each one drives, which of the two was on screen and whether the editor was open in the visualiser. It is saved in presets (loading a preset restores both patterns exactly; loading a preset with no pattern clears the current ones) and in the DAW project (close the project mid-edit and it reopens precisely as you left it). Any pattern edit marks the preset modified, as with every other sound edit; opening or closing the sequencer view alone does not and neither does switching between the two sequencers.

11.7.9 MUTING AND PAINTING

Muting steps - click the step number at the top of a column (lit = active, dimmed = muted). Each lane's mutes are independent: mute on GATE silences the note entirely (ratchets included); mute on PITCH plays the held note with no offset; mute on VEL falls back to your played velocity; mute on an SM lane outputs zero. Dragging a gate slider to 0% is not a mute - a near-zero gate still fires the note and the envelope's release will ring on.

Painting values - hold the mouse button in any lane's slider area and drag across the columns: every step you cross takes the value at the cursor's height, so you can draw curves, ramps and shapes in a single gesture.

11.7.10 ONE CLOCK

Everything tempo-synced in the Instrument - sequencer steps, tempo-synced LFOs, envelope Sync retriggers and grain Sync boundaries - derives from a single musical grid: the DAW's song position while the transport runs and a shared free-running clock at the host tempo while it doesn't. Boundaries that fall on common multiples are coincident to the sample: a 1/16 sequencer step, a 1/4-synced LFO cycle and a half-bar envelope retrigger all land together, in both transport states.

11.7.11 A FEW QUICK RECIPES (PRE-DEFINED)

Classic acid line - GATE lane: a mix of short steps and a couple of LGT steps; enable Voice Legato with some glide; PITCH lane Q 12 with a few +12/-12 jumps; RATE 1/16; SWING 55-66%.

Polyrhythmic motion - unlink lanes; GATE 16 steps, PITCH 12, VEL 8; direction Forward. The pattern evolves over 48 steps before repeating.

Rolls and fills - set x3 or x4 ratchets on the last step or two of the GATE lane; drop VEL slightly on the ratcheted steps for a natural fill.

Timeline-locked groove - RETRIGGER OFF so every key lands in step with the arrangement no matter when you play it.

11.7.12 TWO SEQUENCERS AND OSCILLATOR ROUTING

The Instrument has two step sequencers, called sequencer 1 and sequencer 2. They are edited in the same panel, one at a time. The SEQ 1 / SEQ 2 chip in the top row tells you which one you are looking at and clicking it switches to the other. Everything in the panel follows the switch: the lane tabs, the steps, the playhead, the right-hand column and the LOAD and SAVE buttons, which always act on the sequencer you can see.

Sequencer 2 is a complete sequencer in its own right. It has its own gate, pitch and velocity lanes, its own rate, swing, gate scale, direction, retrigger, latch and arp settings and its own enable dot. Because each sequencer keeps its own rate and its own lane lengths, the two can run in different time against each other: sequencer 1 at 1/8 with a 13-step gate lane against sequencer 2 at 1/16 with an 8-step lane is a polyrhythm that takes a long time to come round. The one thing sequencer 2 does not have is the four SM modulation lanes. Those belong to sequencer 1 only, so there is exactly one SM1, one SM2, one SM3 and one SM4 in the modulation source list and nothing about modulation routing changes whether you use one sequencer or two.

Both sequencers are driven by the same keys. Hold a note and every sequencer that is enabled plays its pattern from that note. What decides which oscillators you actually hear each sequencer on is the row of three small squares in the top right of the panel, labelled 1, 2 and 3 for the three oscillators. Each sequencer has its own row. A lit square means that sequencer drives that oscillator; a dark square means it does not. So with sequencer 1 showing square 1 lit and squares 2 and 3 dark and sequencer 2 showing square 1 dark and squares 2 and 3 lit, oscillator 1 plays sequencer 1's pattern while oscillators 2 and 3 play sequencer 2's.

The squares follow one rule that keeps the two sequencers from fighting over an oscillator. Turning a square on claims that oscillator for the sequencer you are looking at and turns the same square off in the other sequencer and in Starling's row too (Section 11.8.5). Turning a square off does nothing anywhere else. That second half matters: because turning a square off never turns anything on, you can leave an oscillator with no square lit anywhere and an oscillator that no sequencer has claimed simply plays the keys you hold, directly, exactly as it would with both sequencers switched off. In the example above, if you also turned square 3 off in sequencer 2, oscillator 3 would play your held notes under the two sequenced parts.


One practical note. The two sequencers share the synthesiser's voice pool with everything else you play and UNISON multiplies what every note costs, so two sequencers with long gates on a high UNISON setting will run out of voices before anything else does. If sequenced notes start cutting short, lower UNISON or raise the Voice Limit (Section 11.6).

11.7.13 COPYING A LANE

Any lane can be copied to any other lane, on the same sequencer or the other one. Right-click a lane tab - GATE, PITCH, VEL or SM1 to SM4 - and a small menu offers Copy and Paste for that lane. Copy takes the whole lane: every step's on/off, value, legato, octave, ratchet and chance, together with the lane's length, Q and mute. Paste is only offered when there is a lane on the clipboard. To move a pitch line from sequencer 1 to sequencer 2, right-click PITCH and choose Copy, click the SEQ chip, right-click PITCH again and choose Paste.

A lane does not have to be pasted onto a lane of the same kind. SM1 pasted onto GATE or GATE onto VEL, carries the step values across as they are, which is a quick way to turn a rhythm into a modulation shape or the other way round. If LINK LANES is on in the sequencer you paste into, the pasted lane's length is applied to every lane, exactly as dragging the length slider would. The lane travels on the system clipboard, so it can also be pasted into another instance of NeuralDream.

11.8 - STARLING



"Starling listens to your playing and to its own audio; captures a theme without ceremony; hears chords as chords and finds your progression by its repeat; detects key and mode from function rather than note-counting; plays on the DAW's grid with your own feel; wanders and returns on a shaped arc; judges its own phrases before you hear them; takes up its own material when you stop and comes home when it's been away too long; voices chords under its melody; drives four modulation lanes; and draws its past, present and future."

11.8.1 WHAT STARLING IS

Starling is a generative musical engine that lives inside the Instrument, alongside the step sequencer. It is not an arpeggiator and it is not a pattern player. You play; it listens; it takes what you played as a subject and develops it - departing from it, returning to it, accompanying you while your hands are on the keys and taking the lead when they are not.

The step sequencer plays what you drew. Starling plays what it heard.

There are many MIDI generators, sequencers and arpeggiators in the world. What makes Starling different is that it listens to the audio it is causing. It knows not only which notes it generated but how they actually sounded on this patch, through this chain of engines, in this room of reverb - whether a note landed or was lost under the previous tail, whether the texture is already dense enough. That feedback is only possible because Starling lives inside the synthesiser rather than in front of it and it is the reason Starling behaves like a player listening to a room rather than a sequence being read from a page.

11.8.2 WHY "STARLING"

A murmuration of starlings has no leader and no choreography. Each bird follows a handful of simple rules about the birds nearest it and out of that comes a shape that turns, folds and reforms - coherent, purposeful and impossible to predict in detail. It looks composed but nothing composed it.

That is exactly the behaviour we wanted: a musical intelligence with rules rather than a score, producing something that has direction and shape without ever being pre-ordained. The name carries the design. It also carries the display, which draws the sounding notes as a small flock holding station and drifting - and the whole engine's guiding principle, which is that the flock stays together. Starling always has a home to return to and that home is what you played.

11.8.3 OPENING STARLING

Click STARLING in the preset strip, to the right of the SEQUENCER button. The main visualiser switches to the Starling view; click again to return. If the visualiser is folded away (see Section 8.2), it opens and XFM switches on; folding the visualiser away again closes Starling. SEQUENCER and STARLING are two views of the same space - click either to switch straight to it.

The view has two halves. On the left is the field: a circle with the twelve pitches arranged around it, where Starling draws what it is doing. On the right is a column of controls.

Starling always loads switched off. A project saved with Starling running will not start generating the moment you open it, which would be startling at best and unpleasant if the monitors have been turned up since. Everything else about Starling - your theme, its key, its settings - is restored exactly as you left it, so switching it on resumes where you were.

11.8.4 GETTING STARTED

Switch STARLING on with the copper dot at the top of the column and play.

That is the whole procedure. There is no arming, no record button and no capture step. While your hands are moving Starling listens; the moment you pause, whatever you just played becomes its subject and it answers you in that gap. The field tells you what it is waiting for before you begin.

If you play a single line, a breath's pause is enough. If you play chords, it waits for a full bar of silence before deciding you have finished - because the gap between two held chords is about the same length as a short rest and it must not mistake one for the other. If you play a chord sequence twice, Starling recognises the repeat and takes the whole progression as your subject without waiting for you to stop at all.

CLEAR THEME - forgets the subject; the next thing you play becomes the new one.

11.8.5 THE CONTROLS

Each control governs a different axis. None of them overlaps another.

DRIFT - how far each variation departs from the subject. At low settings you hear your phrase with a note moved; at high settings you hear it developed - a half restated a step higher, a shape turned upside down, a call answered rather than repeated. DRIFT decides distance, not direction.

GRAVITY - how often Starling comes home. It sets the length of the trip out and back and the row reads in bars, not as a percentage, because the number that matters is how long you wait to hear your theme again. A four-bar theme with an eight-phrase cycle is thirty-two bars between statements; the display says so. High gravity gives short trips and frequent returns; low gravity gives long arcs.

SELF - whose material becomes the next subject. At 0% Starling only ever develops what *you* played. Turn it up and, once you have stopped for a while, it begins taking up its own phrases as the thing it develops, so the music moves on instead of circling one idea. Play again and your new phrase becomes the subject immediately. Starling never falls silent while it decides - it plays continuously throughout; SELF governs only which material seeds what comes next. Your original theme is held separately and is never overwritten by this: after a couple of cycles on its own material Starling returns to what you played.

HARMONY - how strongly strong beats are pulled onto the chord Starling hears under each bar. At 0 only the key and scale apply and any note in the scale may fall anywhere. Turn it up and downbeats land on chord tones while the notes between them stay free to pass through. It affects where notes land, never the shape of the phrase.

CHORDS - how much Starling voices chords beneath its own melody. At 0 it plays a single line. Turn it up and chord tones appear under the strong beats - a single supporting voice at lower settings, a full triad at the top - always below the melody, always softer and always thinning themselves while you are playing, because two people voicing the same chord is mud.

LISTEN - the ear. This is the control that makes Starling different from a MIDI generator: it governs how much the *actual sound* of your patch throttles what Starling plays. On a long, washy preset, the reverb tail is already filling the space, so notes are shortened, softened and sometimes left out entirely. On a dry, short sound the same setting barely changes anything, because there is nothing in the way. LISTEN affects density, never direction.

KEY and SCALE - the note collection. Both default to AUTO, in which Starling works out the key and the mode from what you play and displays its verdict - "AUTO - C", "AUTO - Minor". It listens for function rather than counting notes: the notes that begin bars and phrases, the tonic chord sounding as a shape, the note you open on and the note you land on. Choose a key or scale by hand to override it. Set SCALE to Chromatic to leave pitches exactly as played.

METRONOME - a click that follows the DAW's tempo and time signature, with a brighter tone on the first beat of each bar. Starling measures everything you play against beats and bars, so if you cannot hear them - the transport stopped or your DAW's own click switched off - you have nothing to play to and Starling's reading of your phrasing will suffer. This is that reference. It is a global setting rather than part of a preset, so it never starts ticking when a sound loads. Note that the click is part of the plugin's audio output, so switch it off before bouncing.

The click belongs to Starling, so it is only ever heard while Starling is switched on somewhere in the patch. Switching Starling off silences it and it never starts on its own when a project loads. Stopping the transport stops it, as you would expect, but switching the metronome on while the transport is stopped starts it ticking at the project tempo, so you can play along with Starling without the DAW running. Starting the transport again puts it back in step with your DAW.

LOAD and SAVE - save the current subject and its settings as a .star file and load one into any preset. This is how you carry a subject from one sound to another: capture a phrase on a piano, save it, load it into a pad and let Starling continue with the same material in a different voice. The files live alongside your presets, in the DreamDeep/NeuralDream folder inside your user Library. These three controls are available while Starling is stopped and are greyed out while it is running.

11.8.6 THE FIELD

The display draws four things.

The ring places the twelve pitches around a circle of fifths, so neighbouring positions are musically neighbouring keys. A consonant chord draws as a narrow shape; a tense one throws spikes. The centre is home - your theme.

The flock is what is sounding now: each note a bird holding its station on the ring and drifting gently. As Starling departs from your theme the flock moves away from the centre; as it returns, the flock comes home.

The trail is the copper path behind the flock - where it has actually been over the last few seconds.

The plan is the cool blue dashed path ahead of it. Starling is the only engine in the plugin that knows its own future: a phrase is entirely determined by its position in the cycle, so the next eight phrases can be drawn before they are played. Where the trail and the plan diverge, you are seeing Starling reject a phrase it had planned and choose a better one.

Along the bottom, the phrase strip shows the trip out and back - one cell per phrase, the current one lit, the cells at each end being home - with a caption naming the movement, the phrase, the chord under the current bar and what Starling is doing: THE THEME, STATED, WANDERING 40% or ACCOMPANYING YOU.

11.8.7 MOVEMENTS

Over a long passage, a piece changes character - it moves from something quiet to something energetic and back. Starling follows this. When your playing shifts in energy or register and *stays* shifted for several bars, Starling declares a new movement and re-anchors itself to the new character: the material of a slow section is never dragged into an energetic one. Your original theme still returns but as a deliberate recapitulation rather than by accident. The caption shows which movement you are in.

11.8.8 ACCOMPANYING AND LEADING

While your hands are on the keys, Starling thins and softens itself - it frames what you are playing rather than competing with it and the caption reads ACCOMPANYING YOU. Stop for a couple of beats and it takes the lead, filling out to its full density. This needs no setting; it is simply how Starling behaves.

11.8.9 S1-S4 - STARLING MODULATION SOURCES

Starling generates four modulation sources, available anywhere in the plugin that accepts one, listed immediately after SM1-SM4:

Source	Carries
Star Note	The pitch of the note currently sounding
Star Velocity	How hard that note was struck
Star Shape	The slow amplitude arc of the passage - its crescendo and diminuendo
Star Gate	The length of the current note

These are Starling's counterpart to the sequencer's SM1-SM4, with one important difference: where you *draw* the SM lanes, Starling *generates* these. Route Star Velocity to filter cutoff and the filter opens with the force of the playing; route Star Gate to an envelope decay and long notes breathe differently from short ones; route Star Shape to volume or resonance and the whole patch swells and fades with the music's own arc.

Two behaviours are worth knowing. They are always available - they work whether or not Starling is making a sound and they respond to your playing as well as to Starling's, so you can use Starling purely as a modulation source with its own notes silent. And they hold their last value between notes rather than dropping to zero, so nothing lurches at a rest.

The four small meters in the strip directly above the keyboard, on the right of the field, show what each lane is doing, so you can see a source before you route it.

11.8.10 A FEW STARTING POINTS


Let it answer you. Switch Starling on with everything else at its default. Play a short phrase and stop. It will answer in your first pause. Raise DRIFT to hear it develop the phrase rather than restate it.

Play together. With Starling running, keep playing. It will thin itself and frame you; stop and it takes over. This is the mode in which it feels most like a second player.

Let it go somewhere. Set SELF to around 60% and stop playing. It will develop its own material for a while and then return to your theme - a genuinely different piece of music each time, still recognisably built from what you gave it.

Use it as a modulator. Switch CHORDS and the note output down, route Star Shape to filter cutoff and let Starling breathe the patch while you play over the top.

Two practical notes. Play to the METRONOME or to your DAW's click, so that Starling's reading of your bars matches your intention. And watch GRAVITY: it reads in bars and if your theme is four bars long, a low setting means a very long wait before you hear it again.



12. PARTICLES

12.1 WHAT IT IS

A particle system in film and video does not model one complicated object. It has an emitter that produces a large number of very short-lived elements, each one born with its own size, colour and speed drawn from a spread of values rather than a fixed setting, each changing over its own lifetime and each dying, sometimes leaving smaller elements behind. Smoke, sparks, rain and fireworks are all made this way.

PARTICLES does the same thing with sound. The emitter gives birth to short tones. Each tone has its own pitch, its own wavetable and its own starting position within it, its own length, its own loudness and its own place in the stereo field, all drawn from the ranges you set rather than from a single value. Each one drifts as it sounds, carries its own two envelopes and its own two low-frequency oscillators and dies when its time is up, optionally leaving children behind at a related pitch.

Turn the pitch range wide and the rate high and you have a shimmering noise. Narrow the pitch to a scale and slow the rate and you have a music box. Everything worth having is somewhere between those two and a single control, PITCH MODE, takes you from one to the other.

Nothing in the engine uses the synthesiser's voices, so a dense cloud can never starve a note you are playing. It has its own pool of up to 1024 particles and when that pool is full the oldest particle fades out over a few milliseconds to make room for the new one, so an emitter set too fast degrades gracefully instead of stalling the audio.

12.2 THE PANEL

Selecting Particles as a layer's source opens a panel the same size as the synthesiser's, divided into seven regions whose names are printed in the background: EMITTER, PARTICLE, PITCH, FLIGHT, CHILDREN, MODULATION and VISUALISER.

Every control on the panel, buttons and dropdowns included, can be automated from your DAW, where each one appears under the layer's number followed by Particles and its name, for example L1 Particles Rate. Every knob can also be a modulation destination for the global LFOs, envelope followers and macros described in Chapter 7. That matters more here than elsewhere: because each particle reads these controls at the moment it is born and every particle is born at a different moment, a fast global Sample and Hold LFO on a control such as WT POS gives every particle a different value rather than moving them all together. A control that another setting has made irrelevant is greyed out and cannot be turned. What is live on the panel is therefore always what you will hear.



The **PRESETS** button, at the top right of the PARTICLE region opposite its title, opens a list of Particles settings to load. INIT at the top returns every Particles control to its default. Below it are the Particles settings from every saved preset that has a Particles layer, listed by the preset's name, with the layer number added when a preset has more than one. Choosing one changes only this layer's Particles controls and its wavetable. The rest of your patch, including its effects and modulation, stays as it is. The last one you chose is ticked. A preset you have just saved appears in the list straight away.

Every knob on this panel also moves at a quarter speed while you hold the **Shift** key, for setting a value precisely.

12.2.1 EMITTER

RATE is how many particles are born each second, from a tenth of one to two hundred. **RATE VAR** scatters the birth times: at zero the particles arrive on a perfect grid, which is a legitimate sound and rarely the one you want and at full they arrive quite unevenly while still averaging the rate you set. **VELOCITY** decides how much the velocity of the note you play sets the loudness particles are born with. At 0%, the default, every particle is born at full **LEVEL** however hard you play. At 100% the whole range follows your playing, from silent for the softest note to full for the hardest. In between it scales: at 50% particles are born between half and full level. The velocity of the most recent note sets every particle born after it, until the next note arrives.


BURST throws a whole handful of particles at once, for a single gesture rather than a stream and **BURST SIZE** is how many it throws. It can also be fired by DAW automation to land a burst at an exact point in your arrangement.

ACTIVE gates the emitter. Switching it off stops new particles being born and particles already sounding finish their lives rather than being cut off. **MAX** is the ceiling on how many may be alive at once, from a single particle up to 1024. At one the engine becomes a monophonic instrument.

SYNC locks the emission to your DAW's tempo and **DIV** chooses the division: with **SYNC** on, one particle is born every division of the bar, on the same list of divisions used everywhere else in NeuralDream, from eight bars down to a thirty-second triplet. With **SYNC** on, **RATE** is greyed out because **DIV** has taken its place. With **SYNC** off, **DIV** is greyed out.

When **SYNC** is on and the transport is rolling, the emitter follows the host's own bar position rather than counting from when the plugin happened to start and each particle is placed on the exact sample its beat falls on. A particle emitter on eighths and a step sequencer on eighths therefore strike together rather than alternating.

The **S** and **M** buttons in the top right corner of this region are solo and mute for the emitter's own particles, working exactly like the ones on the synthesiser's mixer. They silence the sound only: the particles are still born, still live and still spawn children, so you can listen to the children alone without changing anything about how the patch behaves.



12.2.2 PARTICLE

WAVETABLE is the same browser the synthesiser's oscillators use and it behaves identically. **RANDOM** sits at the top and gives every particle a different table, which is what makes a cloud sound like many things at once. Below it is a **SIMPLE** folder of eighteen single-cycle shapes and below that the wavetable library's own folders. The first eight simple shapes (Sine, Saw, Ramp, Triangle, Square, Pulse, Trapezoid and Parabolic) are exactly the synthesiser's own, played by the same code. The other ten, from 2 Harmonics to Metallic, are each built from a handful of harmonics. A library table is loaded the first time you choose it and until it is ready its particles play a plain sine for a moment rather than falling silent.

WT POS is which frame of the wavetable a particle starts on and **WT POS VAR** scatters that so no two particles read the same place. On the eight simple shapes from Sine to Parabolic, **WT POS** sets the width of the shape instead, exactly as **POS** does on the synthesiser's oscillators (see Width of the simple shapes in Section 11.3.2.1). At 0% the shape is complete and turning **WT POS** up squeezes it narrower until it is silent at 100%. Each particle takes its own width and **WT POS VAR** gives every particle a different one. The ten harmonic shapes have a single frame and no width, which is why the two controls are greyed out on those. The wave display draws the chosen simple shape at its current width.

LIFE is how long a particle lives, from five milliseconds to four seconds. Short is sparks; long is a drone made of many voices. **LIFE VAR** scatters the lifetime and it is what stops a cloud pulsing.

LEVEL is the loudness a particle is born with, before its envelope shapes it and **LEVEL VAR** scatters it so particles sit at different depths rather than all at the front. The engine also balances the cloud for you. The more particles are sounding at once, the quieter each one plays. A dense cloud is then as loud as a single particle. A higher **RATE**, a longer **LIFE** or more children therefore make the cloud thicker rather than louder. A particle that is fading in or out counts for less than one at full level and the balance glides rather than jumps. The cloud therefore does not pump as particles come and go. **LEVEL** sets how loud the whole cloud is.

12.2.3 PITCH

This is where the engine goes from noise to music and **MODE** is the control that takes you there. **RANDOM** draws any pitch within the range, which at speed is coloured noise. **SCALE** draws only the notes lit on the keyboard below, exactly as they are lit, whatever you play. **KEY** draws the same lit notes moved onto the note you play. A C minor keyboard then gives D minor when you play a D. **CHORD** draws only the notes you are holding down on your MIDI keyboard, at the octaves you are holding them and falls back to **SCALE** when you are holding nothing. **HARMONIC** draws from the harmonic series of the root, which sounds like something being shattered rather than played. The controls a mode does not use are greyed out. **RANDOM** and **CHORD** grey the keyboard, **SCALE** and **ROOT**. **HARMONIC** greys the keyboard and **SCALE** but keeps **ROOT**, because the series is built on it.

The keyboard is one octave of twelve toggles, C to B, exactly like the one under the **GRAINS** pitch quantiser in the synthesiser. A note that is switched off is never emitted. Under each note is a small blue lamp that lights for a moment whenever a particle is born on that note, at any octave, so you can see the cloud playing.

SCALE sets all twelve toggles in one action from a list of thirteen scales, rotated onto the keyboard by **ROOT**: choose Natural Minor with the root on D and the keyboard lights D, E, F, G, A, B flat and C. Changing the root re-rotates whichever scale you chose.

CENTRE is the middle of the pitch range and reads as a note name. **RANGE** is how far either side of it particles may be born, up to five octaves.

TRACK turns the engine into something you play. With it on, two things follow the note you press: the pitch centre moves to that note and the emitter runs only while a note is held, stopping the moment you let the last one go. The sustain pedal counts as holding: with the pedal down, every key you have played stays held until you lift it. Particles already sounding finish their lives, so releasing a key never clicks. **TRACK** does not change which notes are chosen, because that is the job of **MODE**. **KEY** is the mode that moves the scale with the note you play. With

TRACK on the CENTRE knob is greyed out, because the note you play is setting it. With TRACK off the emitter runs all the time, needs no MIDI at all and CENTRE sets the middle of the range.

12.2.4 FLIGHT

What a particle does between being born and dying. PITCH DRIFT is how far its pitch moves over its life, in semitones, negative to fall and positive to rise and DRIFT VAR scatters that so they do not all fall together.

PAN is where a particle is born in the stereo field, from minus 100 hard left through zero in the centre to plus 100 hard right. PAN SPREAD is how far either side of that position particles may land, so for a genuinely hard left cloud you need PAN at minus 100 and PAN SPREAD at zero. PAN DRIFT is how far a particle travels across the field while it sounds and WT DRIFT how far it moves through its wavetable's frames, which on a library table is a timbre sweep over its life.

12.2.5 MODULATION

These four modulators belong to the particle rather than to the plugin and the reason is worth understanding. A global LFO moves every particle in exactly the same way at the same moment, so the whole cloud swerves together. To make particles move differently from one another while they sound, each one has to carry its own modulators, struck at the instant it is born.

ENV 1 is a four-stage envelope wired to the particle's level. Every particle needs one, so it is not optional. With DECAY at zero the decay fills whatever time is left after the attack and the release, which means the release always shapes the end of the particle. There is no note-off here: a particle's release is simply the last few moments of its life.

ENV 2 is a second four-stage envelope which you send wherever you like using DEST, choosing pitch, wavetable position, pan or level, with AMOUNT setting how much of it arrives. Pitch is its obvious first job. Neither envelope loops or syncs, because a particle is usually far too short for either to mean anything.

LFO 1 and LFO 2 each have a rate, a shape, a destination and an amount and can be locked to the DAW tempo with their own SYNC and DIV. With an LFO's SYNC on its RATE is greyed out. With its SYNC off its DIV is greyed out. The shape list includes Sample and Hold, which matters more than it sounds. PHASE decides where each particle's LFO starts: RANDOM gives every particle a different starting phase so the cloud shimmers, ZERO starts them all together for a deliberate synchronised gesture and FOLLOW takes the phase from the matching global LFO at the moment of birth. RATE VAR scatters the rate so particles do not all wobble at the same speed. Sample and Hold with a random phase gives each particle its own stepped wander, which is what turbulence means here.

Independently of all of this, every particle fades in and out over a few milliseconds or a fifth of its life if it is shorter than that. This cannot be switched off and no envelope setting defeats it. A particle system starts and stops hundreds of sounds a second and every one of them would click otherwise.

12.2.6 CHILDREN

A particle that spawns more particles at the moment it dies, which is what turns a spark into a shower. ON enables it and COUNT is how many each parent leaves behind. With ON switched off, every other control in this region is greyed out.

PITCH OFF is where the children sit relative to the parent's pitch at the moment it died, in semitones and OFF VAR scatters that. LEVEL and LIFE are the children's loudness and lifetime as a proportion of the parent's. SPREAD is how widely they scatter from where the parent died.


GEN decides whether children may themselves have children. Two generations is the ceiling on purpose: three is an explosion the pool cannot survive and nobody can steer. Children inherit the parent's wavetable. This region has its own S and M buttons, so you can hear the shower without the shells or the shells without the shower.

12.2.7 THE TWO DISPLAYS

The VISUALISER region on the panel shows the wavetable the particles are reading. Every table is drawn by exactly the same code as the synthesiser's wavetable display, so the two look and behave alike: every frame of the table is drawn as one thin line, brightest at the front and fading into the distance, over a backdrop running from black at the top to deep blue at the bottom. A library table shows each of its frames. A simple shape, from Sine to Parabolic, shows its widths instead, from WT POS 0% at the back to 100% at the front. A harmonic shape has a single cycle and is shown as a stack of that cycle, as the synthesiser shows a single cycle. RANDOM has no single table to draw and the display says so in words. Here the lines are drawn in light blue and the frame WT POS is reading is drawn in white on top so you can follow it. The white frame follows WT POS as it is actually playing, including any modulation routed to it. An LFO on WT POS therefore sends it travelling back and forth through the table. It shows WT POS itself, while WT POS VAR, WT DRIFT and the particle LFOs scatter individual particles away from it. Drag inside it to turn the landscape, up and down to tilt it and double-click to put it back to straight on; scrolling sideways slides the view along the cycle and scrolling up and down zooms in and out, as described in Section 11.5. The button in its corner switches between the three-dimensional view, which is the default, and a flat view of the single frame being read. Your choice of view is saved with the preset.

Along the foot it shows the table's name, the last note received, the pitch centre in use and how many particles are alive against MAX, in amber when the pool is full.

The plugin's main visualiser, at the top of the window, shows the cloud itself as a three-dimensional box. Left and right is the particle's position in the stereo field, up is its pitch and depth is its age: a particle is born at the back, where the emitter glows and flies towards you as it sounds, dying at the front. So the rate is how densely the back of the box is populated, the variation controls are how scattered they are, a pitch or pan drift shows as a curved flight path and a child is born exactly where its parent died and flies off again. First-generation particles are cyan, children amber and grandchildren magenta and the camera sways slowly so the depth reads even when nothing is changing.




12.3 EIGHT STARTING POINTS

These are recipes rather than presets, because only your ear can settle the levels. Each one is the engine alone in a single layer and every control not listed is at its default.

Name	Category	The settings that matter	What you should hear
Music Box	Melodic	RATE 6, RATE VAR 40%, LIFE 0.6 s, SCALE Major Pentatonic, ROOT C, CENTRE C5, RANGE 1.5, a bell-like library table, LEVEL VAR 40%	A slow scatter of pentatonic bells, never the same twice
Cloud	Ambient	RATE 60, LIFE 1.8 s, LIFE VAR 60%, SCALE Chromatic, RANGE 3, WAVETABLE RANDOM, PAN SPREAD 100%, LFO 2 Sample and Hold to pitch at 15%	A dense evolving pad made of hundreds of tiny voices
Rain	Textural	RATE 120, LIFE 0.05 s, LIFE VAR 80%, MODE Random, CENTRE C6, RANGE 2, WAVETABLE Sine, PAN SPREAD 100%, LEVEL 40%	Fine, high, unpitched rain
Sparks	FX	RATE 20, RATE VAR 80%, LIFE 0.12 s, MODE Harmonic, ROOT C, PITCH DRIFT -24, DRIFT VAR 50%, WAVETABLE Glass	Bright shards that fall as they die
Fireworks	FX	RATE 2, LIFE 0.4 s, MODE Scale Major, CENTRE C4, CHILDREN on, COUNT 8, PITCH OFF +12, OFF VAR 80%, SPREAD 100%, GEN 2 gens, child LIFE 80%	Slow shells that burst into showers of showers
Shatter	FX	RATE 0.5, BURST SIZE 128, MODE Harmonic, RANGE 4, LIFE 0.3 s, PITCH DRIFT -12, WAVETABLE Metallic, PAN SPREAD 100%	Nothing until you press BURST, then a bell smashed on a stone floor
Follow Me	Melodic	RATE 30, MODE Chord, TRACK on, LIFE 0.25 s, WAVETABLE Odd Harmonics, LFO 1 Sine to pan at 60%, 4 Hz	A cloud that plays only the chord you are holding, swirling around you
Pulse	Rhythmic	SYNC on, DIV 1/16, RATE VAR 0%, LIFE 0.08 s, MODE Scale Minor Pentatonic, ROOT A, WAVETABLE Square, LEVEL VAR 0%	A tempo-locked sixteenth-note pattern of random pentatonic pulses

For fun with Stereo Panning, try:

Set RATE to 1 or 2, LIFE to about 3 seconds, PAN to the centre and PAN SPREAD to 0
Set ENV 2 to A 1.5 s, D 0, S 100% and R 0.5 s. Set DEST to Pan and AMOUNT to +100%



APPENDIX A - HOW THE LADDER FILTER IS MADE


The LADDER filter in NeuralDream is a model of the transistor ladder circuit rather than a filter designed to have the same frequency response. The distinction matters, because the sound of the original comes from what the circuit does to a signal that is not small and a response curve says nothing about that.

The model follows Antti Huovilainen's 2004 paper, "*Non-linear digital implementation of the Moog ladder filter*", which derives the equations directly from the circuit. Each of the four stages is a transistor differential pair driving a capacitor and the current through a differential pair follows a hyperbolic tangent of the voltage across it. So each stage is computed as a one-pole filter whose input and whose own state both pass through that tangent curve before they meet and the resonance is created, as in the hardware, by feeding the output back into the first stage's pair, where it saturates along with the signal. The saturation is therefore not an effect applied to the filter; it is inside the filter, in five places and it shapes the sound at every level rather than only when driven.

Three practical measures from the paper are used. The stages run at twice the sample rate, because a nonlinearity generates harmonics above the audible band which would otherwise fold back down as aliasing. The feedback carries a half-sample delay, formed by averaging two successive outputs, which keeps the resonant peak on the cutoff frequency rather than letting it drift above it. And because each stage reads the saturated output of the previous one, that value is computed once and reused, so the model needs five saturation evaluations per step rather than eight.

The transistor ladder as NeuralDream computes it

Huovilainen, DAFx 2004. Every saturation shown is a real transistor pair in the original circuit. Everything inside the dashed box runs at twice the sample rate.




The band pass and high pass are not separate filters: they mix the four stage outputs $y1..y4$ with binomial weights.

HP = $x - 4y1 + 6y2 - 4y3 + y4$ BP = $4y2 - 8y3 + 4y4$

Inside one stage

$$y[n] = y[n-1] + g \cdot 2Vt \cdot (\tanh(in / 2Vt) - \tanh(y[n-1] / 2Vt))$$

g is the one-pole coefficient set by cutoff, $1 - \exp(-2 \pi f_c / 2fs)$. $2Vt$ is the transistors' thermal voltage and sets how hot the signal sits against the curve; it is 0.5 in NeuralDream, set by measurement.



NeuralDream v2.15.70. Both saturations in each stage, and the saturation of the feedback, are the transistor pairs of the original circuit; nothing here is an effect added afterwards.

Two refinements are NeuralDream's own. Real transistor pairs are never perfectly matched and an ideal model produces only odd harmonics, which is cleaner than any real ladder. A small standing asymmetry in the input pair's curve restores the even harmonics that give the sound its warmth; it is applied to the low pass, where that character belongs and not to the band and high pass, whose outputs depend on the stages matching. And the feedback needed for self-oscillation rises slightly with cutoff in any discrete model, so it is compensated by a fitted curve, measured at eight cutoffs across the range, so that the RESONANCE control means the same thing at every frequency.

The level at which the signal meets the transistors was set by measurement against what the synthesiser's own oscillators actually deliver to the filter, not by theory: at one oscillator on a normal level the ladder adds around three percent harmonics with the second and third both present, at two it adds around eleven and with everything up it compresses the way the hardware does with everything up. The passband level at ordinary settings is unchanged from the previous filter to within a tenth of a decibel, so existing patches keep their balance and gain only their new character.

One consequence of building the resonance the way the circuit builds it is that the low end thins as RESONANCE rises. The resonance is made by feeding the output back into the first stage and that feedback subtracts from the signal arriving there: the more feedback, the less of the original gets through. It shows up first in the bass, because the bass passes through the filter almost untouched and so meets the full force of the subtraction. With the cutoff at 1 kHz, our ladder measures 0.1 dB down at 50 Hz with the resonance off, 6.8 dB down at a quarter, 10.5 dB down at half and 13.1 dB down at three-quarters. This is not a fault in the model; it is what every 24 dB per octave resonant filter of this design does, in hardware as much as in software and it is why a Minimooog bass line loses its weight as the *Emphasis* goes up.

The **BASS** button cures it, in the same way the analogue designs that offer this do: the signal going into the ladder is lifted by exactly the amount the feedback takes away, which is one plus the feedback. The low end then stays where it was at every resonance setting. On the same measurement, 50 Hz reads 0.1 dB down at a quarter resonance, at half and at three-quarters, instead of 6.8, 10.5 and 13.1. Nothing else about the circuit changes, so the resonant peak still rises as it should and the filter still saturates as it did; the sound is simply fuller and a little louder underneath. The lift is applied only while the filter is a low-pass. On a high-pass or a band-pass the lift would only make the filter louder, because no bass passes through for the feedback to take away.

.

APPENDIX 3 - COPYRIGHT + LEGAL

3.1 COPYRIGHT AND LICENCE

NeuralDream is copyright © DreamDeep Plugins 2026. All rights reserved.

This software is licenced for use by the original purchaser only. A licence allows you to activate NeuralDream, both the effect and the instrument, on up to three computers that you own and use yourself. A trial licence gives the same use for the length of the trial only. You may not copy, redistribute, share, resell, rent, lease or sublicense the software or any of its components to any other person or organisation. You may not reverse-engineer, decompile or disassemble the software, except to the extent that the law allows this regardless of any licence terms.

3.2 TRADEMARKS

NeuralDream is an independent product. Other product and company names mentioned in this manual are trademarks of their respective owners. They are used only to acknowledge the instruments and ideas that inspired NeuralDream and they do not imply any connection with or endorsement by, those companies. VST is a trademark of Steinberg Media Technologies GmbH, registered in Europe and other countries. Audio Units is a trademark of Apple Inc. AAX and Pro Tools are trademarks of Avid Technology, Inc.

3.3 PRESETS AND IMPULSE RESPONSES

The factory presets and impulse responses included with NeuralDream are provided for your use within the plugin and you may use them in your own music productions without restriction. You may not redistribute the preset files or the impulse response files separately from the plugin.

3.4 THIRD-PARTY COMPONENTS

NeuralDream includes the following third-party software, each used under the terms of its licence:

- ◆ **NeuralAmpModelerCore** by Steven Atkinson is used under the MIT licence, for the Neural Amp engine and the neural mode of Tape Simulation.
- ◆ **smbPitchShift** by Stephan M. Bernsee is used under the WOL (Wide Open Licence), for the Pitch Shifter engine and time-stretch processing.
- ◆ **LibTorch**, the C++ distribution of PyTorch, is used under the PyTorch licence, which is BSD-style. It runs the neural network models in the RAVE engine. PyTorch is copyright © 2016 onwards Facebook, Inc. and its contributors.
- ◆ **The VST 3 SDK** by Steinberg Media Technologies GmbH is used under the MIT licence, for the VST3 plugin format.
- ◆ **TweetNaCl** by Daniel J. Bernstein and others is in the public domain and it is used to verify licence files.
- ◆ **The JUCE framework** by Raw Material Software Limited is used under licence, for the plugin framework, the audio processing and the user interface.

The full licence text of each component above is included in the file Third-Party Licences, which is installed alongside NeuralDream.

The neural models used by the Neural Amp, RAVE and Tape Simulation engines are not part of NeuralDream. You download them yourself from their publishers and each model is covered by its own publisher's licence terms.

3.5 WARRANTY AND LIABILITY

NeuralDream is provided "as is" and DreamDeep Plugins gives no warranty, express or implied, beyond those that the law requires. To the fullest extent the law allows, DreamDeep Plugins is not liable for any damages arising from the use of this software. Nothing in this licence affects your statutory rights as a consumer and nothing in it limits liability that cannot lawfully be limited.

3.6 SUPPORT

For support, updates and more information, visit dreamdeepcr.com.