



VERDALIS
Natural sound synthesis

WhooshPact

Synthesised transitions and impacts

Version 0.4.0

CLAP instrument for Linux and Windows

Part of the Verdalis Plugin Suite

2026-10-04

Contents

Introduction

- Why not just use samples
- The one number the design rests on
- What these references are, and why it matters
- What is in this manual

Installing

- Linux
- Windows
- Hosts
- Your own presets
- Building from source

A first sound

- Notes are triggers, not gates
- Three things to try first

How it works

- The gesture
- Type is a spectrum, not a mode
- Four layers, and two ways of being driven
- Flutter, and what the measurement actually said
- Nothing above 25 Hz is thrown away
- The EQ
- Variation

The plugin window

- Gestures
- The mixer

Parameter reference

- Gesture
- Air
- Tone
- Motion
- Sub
- Hit
- Flutter
- EQ
- Envelope
- Output

The preset library

Using it in a host

- Notes
- Automation and modulation
- Reproducible renders
- CPU cost
- Rendering and tails
- Headroom

What is still being fitted

Reference and limits

- Where the model comes from
- What it does not do

Introduction

WhooshPact is a CLAP instrument that makes the sounds a track is glued together with: whooshes, risers, braams, booms, impacts, accents and downshifters. It contains no samples. Every gesture is built from noise, oscillators, filters and envelopes while the plugin plays.

And no two are alike, which is the point of it.

Why not just use samples

Because a sample plays the same transition every time. Four whooshes in a row out of one file is one sound repeated four times, and the usual way round it is to keep twenty variants of the same whoosh on disk and remember which one went where.

Here every note draws its own span, peak position, filter cutoff, sweep, pitch, level, pan, flutter rate and transient tone from a distribution around the settings on screen. *Variation* decides how far — and setting it to zero makes the plugin bit-exactly reproducible, which is what a transition that has to land on a picture cut needs.

The one number the design rests on

The reference library is 211 production sounds filed in six families. *Peak* — where the loudest moment sits, as a fraction of the whole gesture — separates them more sharply than anything else that was measured:

FAMILY	PEAK SITS AT	SPAN	<100 HZ	CENTROID SWEEP
booms	2.4 %	5.50 s	0.97	−0.12 oct
downshifters	2.4 %	4.57 s	0.98	−1.07 oct
accents	4.3 %	3.55 s	0.69	−0.69 oct
impacts	6.0 %	4.34 s	0.92	−0.63 oct
braams	14.8 %	5.63 s	0.74	+0.60 oct
transitions	32.8 %	3.71 s	0.62	−0.67 oct

A hit peaks in the first few per cent of itself. A whoosh peaks a third of the way in. That is the difference between the two, and it is one knob.

So WhooshPact has no modes. The gesture is five numbers — *Span, Peak, Hold, Rise, Fall* — and those five reproduce every measured envelope contour in the library to within the library's own spread. *Type* selects a family's measured *spectral* profile and nothing else.

What these references are, and why it matters

Every other plugin in the Verdalis suite is fitted against recordings of the world, and the physics is there to be found. These are not recordings of anything: they are finished production sounds, made by somebody else out of synthesis and processing. There is no underlying object to model, and pretending otherwise would mean inventing a mechanism the material does not contain.

So the measurement here describes the *target* rather than the mechanism. That is a real departure from the rest of the suite and it is worth knowing before reading the rest of this manual. The suite's other rule is unchanged: nothing is sampled, and every layer is computed while the plugin plays.

What is in this manual

Installing it, getting a first sound, how the model works and why it is built this way, the window, then a generated reference for every parameter and every factory preset. The last two chapters are honest about what does not fit yet and where the model comes from.

Installing

WhooshPact is one self-contained `.clap` file plus its `presets` folder. There is no installer, no shared library to place and nothing to register.

Linux

CLAP hosts scan `~/ .clap`. Copy the plugin folder from the release archive into it so that you end up with:

```
~/ .clap/WhooshPact/WhooshPact.clap
~/ .clap/WhooshPact/presets/
```

Then rescan plugins in your host. X11 and Cairo are the only external dependencies of the plugin window, and any Linux machine that can run a DAW has both.

The VST3 goes to `~/ .vst3/WhooshPact.vst3`, as a folder, not a file.

Windows

CLAP hosts scan the common CLAP folder. Copy the plugin folder into:

```
C:\Program Files\Common Files\CLAP\WhooshPact\
```

so that `WhooshPact.clap` and `presets\` sit inside it, and rescan. There are no DLLs to install beside it — the window and its Cairo are linked in. The VST3 bundle goes into `C:\Program Files\Common Files\VST3\`.

Keep the plugin and its presets together. The plugin finds its factory presets by looking for a `presets` directory next to its own binary. Moving the `.clap` file on its own leaves it with no factory presets in the *host's* browser. Its own browser still has them: the library is embedded in the binary.

Hosts

Tested with Bitwig Studio and Reaper. In Bitwig, `~/ .clap` is scanned by default; after installing, restart it or rescan under *Settings* → *Locations* → *Plug-in Locations*. WhooshPact then appears as an instrument under the vendor *Ravetracer*, and the factory presets are indexed through CLAP's preset discovery, so they show up in the host's own browser as well as in the plugin window.

Your own presets

`SAVE` in the plugin window writes to a user preset folder, which is created on first use:

PLATFORM	LOCATION
Linux	<code>\$XDG_CONFIG_HOME/WhooshPact/presets</code> , or <code>~/ .config/WhooshPact/presets</code>
Windows	<code>%APPDATA%\WhooshPact\presets</code>

User presets appear in the browser immediately, listed after the factory ones. They are plain text, so they can be copied between machines, kept in version control, or edited in any editor.

COLLECTIONS

A preset library is a shelf, and a long one is worth tidying. Your own presets are filed in **collections**: folders one level under the user preset directory. The browser lists the library down its left side — *All*, the factory presets, each of your collections, and *Unfiled* for presets saved into none — each with the number of presets in it. Clicking one lists just that one.

IN THE BROWSER	WHAT IT DOES
NEW	Makes a collection. It is listed at once, empty until something goes into it
RENAME	Renames the selected collection; its presets go with it
DELETE	Deletes the selected collection and every preset in it, after asking
Drag a preset onto a collection	Moves it there — or onto <i>Unfiled</i> to take it out of one
Right-click one of your presets	Renames it, changes its description, or deletes it

SAVE asks where a preset goes: pick a collection from the list — or *New collection...* to make one on the spot — then give the preset a name and, if you like, a **description**. The description is what the browser shows while the pointer is over the preset and what the help line shows while it is loaded: a sentence for whoever loads it next. Saving over one of your own presets offers its collection and its description again.

The factory presets are locked. They are part of the plugin, not files of yours, so they cannot be moved, renamed or deleted; to change one, load it, adjust it and save it as your own.

Collections are one level deep on purpose: a preset library is a shelf rather than a filesystem, and a tree deep enough to get lost in is one somebody will get lost in.

PRESET PACKS

A whole collection can be written out as a single file — a **preset pack**, extension `.whooshpackpack` — so a library can be handed to somebody else, or moved between machines, as one file rather than a directory.

With a collection selected in the browser:

BUTTON	WHAT IT DOES
EXPORT	Writes the selected collection — its presets and nothing else — as a pack into the <code>packs</code> directory beside <code>presets</code>
EXPORT AS...	The same, to a location you choose
IMPORT...	Lists the packs already in <code>packs</code> , plus <i>Other file...</i> for one from anywhere
REVEAL	Opens the folder the pack was just written to in the system's file browser

Importing adds the pack's presets to the library as a new collection named after the pack. Nothing is ever overwritten: importing the same pack twice gives two collections rather than a mixture of both versions in one.

A pack is the preset format again with a separator line between the presets, so it can be read, diffed and edited by hand like everything else here. It carries each preset's text verbatim rather than a re-serialised copy, which is what makes a round trip through a pack lossless.

`EXPORT AS...`, *Other file...* and `REVEAL` need the desktop's own file chooser — `zenity` or `kdiallog` on Linux, `xdg-open` to reveal, and the operating system's own on Windows. Where none is installed the buttons that need one are not drawn, and the plugin's own `packs` directory, which is where `EXPORT` writes, is enough on its own.

Building from source

The suite builds with a C++17 compiler and CMake 3.16 or newer. From the plugin folder:

```
./install.sh
```

That configures, builds, runs the self-test and installs to `~/.clap/WhooshPact/`. `--vst3` builds and installs a VST3 as well. `WhooshPact_PREFIX=/some/where` `./install.sh` overrides the destination, and `-DCLAP_INCLUDE_DIR=/path/to/clap/include` points the build at a CLAP checkout it cannot find by itself.

A first sound

1. Load WhooshPact on an instrument track.
2. Play middle C. One note is one gesture: it plays out over its own *Span* whether the note is held or not.
3. Open the preset browser by clicking the preset name in the header and try **Simple Whoosh**, then **Deep Impact**, then **Big Horn**.

Those three are the span of the instrument. The first is noise through one filter; the second is a struck low end with a transient on it; the third is a detuned oscillator stack that swells and opens.

Notes are triggers, not gates

This is the one thing about WhooshPact that is not like a normal synthesiser.

The plugin cannot know how long a note will be when it starts, so the note does not decide the length of the sound — *Span* does. A note-on starts a gesture and the gesture plays out.

Letting go early still works: the ADSR runs on top of the gesture, so a short *Release* cuts a gesture short when the note ends. With *Sustain* at 1 and a long *Release*, which is what most of the factory presets have, the note length does not matter at all.

Three things to try first

Move *Peak*. Load *Simple Whoosh* and pull *Peak* down towards zero. The whoosh turns into a hit, because that is the only structural difference between the two. Everything else on the panel stays where it is.

Turn *Variation* up and play the same note five times. At 0.35, which is the default, each one is noticeably its own sound and unmistakably the same preset. At 1.0 they are a family rather than a sound. At 0 they are identical, every time, on every machine.

Turn *Flutter* up on a downshifter. Load *Analog Fall* and move *Start Speed* and *End Speed* apart. The chopping glides exponentially from one to the other across the gesture — 3 Hz at the top to 14 Hz at the bottom in that preset — which is the thing that is genuinely awkward to do with an LFO.

How it works

The gesture

Everything is driven by one clock that runs from 0 to 1 across *Span*. Five numbers shape it:

- *Peak* — where the loudest moment sits.
- *Hold* — how long it stays there.
- *Rise* — the curvature of the approach, as an exponent. 1 is a straight ramp; above 1 holds back and then arrives.
- *Fall* — the curvature of the decay, on the same scale.

The measured transition contour, on a normalised time axis:

t	0.00	0.07	0.13	0.20	0.27	0.33	0.40	0.47	0.53	0.60
env	0.012	0.065	0.162	0.348	0.519	0.560	0.528	0.388	0.299	0.188

That is $(t/0.33)^{1.8}$ on the way up and about $(1-u)^{2.5}$ on the way down. The defaults are exactly those numbers, and the same five fit all six families.

A gesture whose *Span* is longer than what you can hear of it is normal. A boom with *Peak* at 0.02 and *Fall* at 2.2 is below -45 dB long before its span ends; the span is the clock, not the audible length.

Type is a spectrum, not a mode

Type carries two measured numbers per family and nothing else: the slope of the family’s octave-band curve above 125 Hz, and how far 63 Hz stands over 125 Hz. Both are relative to Transition, which is the broadest curve in the library and therefore the neutral setting.

FAMILY	SLOPE ABOVE 125 HZ	63 HZ OVER 125 HZ
accents	-3.73 dB/oct	+7.4 dB
booms	-9.13	+12.5
braams	-5.23	+4.0
downshifters	-10.18	+10.6
impacts	-4.43	+7.0
transitions	-4.08	+2.3

Blend crossfades to the next family, so the six are a continuum rather than six settings.

It deliberately does **not** set the shape in time. If it did, choosing a type would fight the preset that had just set *Peak* and *Hold*, and the five gesture numbers are exactly the ones a user wants to reach for.

One caveat on how the numbers are used: a true -10 dB/octave would take 60 dB off the top of a boom and leave nothing to mix with. The plugin applies 2.2 octaves’ worth of the measured slope through a single high shelf. The measurement is a statement about balance, and that is how it is spent.

Four layers, and two ways of being driven

Air and **Tone** follow the gesture envelope. **Sub** and **Hit** are *struck* at a point in the span and then decay on their own.

That division is not a convenience. A boom is one event with a decay, not a curve with a low end under it, and the measured per-band decays say so: an impact’s 63 Hz band takes 0.78 s to fall 20 dB where its 8 kHz band takes 0.21. Those are two different envelopes on two different layers, and trying to make one gesture curve carry both is what makes a synthesised impact sound like a snare.

AIR

Noise of a chosen colour through one filter, swept across the gesture.

Six colours, each white noise with one filter on it, which is what those names have always meant: White is flat, Pink falls 3 dB/octave, Brown 6, Blue and Violet rise by the same, and Green is a broad mid emphasis around 500 Hz. Each is normalised to the RMS of white, so the *Noise* knob changes colour and not level.

Resonance narrows the filter as well as sharpening it. At 0 it is a plain lowpass and the whoosh is a wall of air closing down; at 1 it is a resonant bandpass and the whoosh has a pitch to follow. The two go together because that is how the ear reads it: a narrow moving band is a pitch, a broad one is wind.

Width is decorrelation. At 1 the two channels are independent noise, which is as wide as a noise source gets; at 0 they are the same signal. The library's transitions measure a median L/R correlation of 0.60, and the default renders at 0.61.

THE SWEEP HAPPENS LATE

A transition's spectral centroid is flat for its first third and then falls 0.80 octaves:

t	0.00	0.07	0.13	0.20	0.27	0.33	0.40	0.47	0.53	0.60
oct	0.00	0.04	0.14	0.02	0.04	-0.05	-0.28	-0.53	-0.61	-0.56

An envelope that moves the filter evenly across the gesture sounds like none of these. *Curve* is where that comes from: at 1 the sweep is spread evenly, above 1 the cutoff sits still and then moves late.

Sweep is in octaves of **cutoff**, and a centroid does not move as far as the cutoff carrying it — what is left below the corner does not move at all. The default of -1.6 octaves of cutoff renders as -0.56 octaves of centroid, against the library's -0.67.

Braams are the one family that rises, by +0.60 octaves, and every braam preset therefore has a positive *Sweep* and a positive *Glide*.

TONE

Three oscillators, detuned against each other, with a pitch glide across the gesture on the same *Curve* warp the filter uses.

Saw and Square are band-limited with polyBLEP. That is not fastidiousness: a braam climbs an octave or more, and a naive saw edge at the top of that climb folds back audibly. Supersaw is the same stack with the detune spread three times wider.

The measured median fundamental of the library's pitched families is 41–49 Hz — booms at 49, braams at 47, downshifters at 41, all within a tone of G1. That is where *Pitch* and *Sub Pitch* default. The played note transposes from middle C, so a preset plays as it was fitted when middle C is played.

SUB

A driven sine with a pitch drop, struck at the gesture's peak.

Drop is in semitones and happens across the layer's own decay rather than across the gesture: a boom's pitch falls as it dies away, not as the whole gesture runs. *Drive* is saturation on the low layer alone — a pure sine at 45 Hz disappears on a small speaker and the harmonics this adds are what carry it there. *Click* is a very short upward pitch snap at the moment of impact, and it is what makes the ear hear something landing rather than a note fading in.

97 % of a boom's energy and 98 % of a downshifter's is below 100 Hz. This layer is not a garnish.

HIT

Three inharmonic resonators at 1, 1.71 and 2.63 times *Tone*, morphing towards plain filtered noise as *Noise* is raised, plus a low *Body* under them. Struck at *Time*, which is a fraction of the span.

Time is where a whoosh-hit comes from: the noise layer peaks two thirds of the way in and the transient is timed to arrive exactly there. The library files fifteen of those under *Whoosh Hit*, and the timing is the whole trick.

Inharmonic ratios on purpose. Struck metal is inharmonic, and a harmonic series here reads as a bell rather than as debris.

Flutter, and what the measurement actually said

The plugin has *Flutter* because chopping a transition up is a staple and doing it with an LFO means drawing an automation curve. But the measurement was worth running before building it, and it produced a mostly negative result that changed the design.

Taking each reference’s amplitude envelope with its slow shape divided out, and asking what fraction of the modulation energy sits at any single rate:

FAMILY	RATE, FIRST THIRD	RATE, LAST THIRD	ENERGY AT ONE RATE
accents	7.7 Hz	7.0 Hz	0.031
booms	9.1 Hz	8.4 Hz	0.022
braams	6.5 Hz	5.6 Hz	0.031
downshifters	8.1 Hz	16.4 Hz	0.045
impacts	6.8 Hz	7.7 Hz	0.038
transitions	6.8 Hz	7.4 Hz	0.033

A metronome would give 1.0 in the last column. **Five of the six families are not fluttering at all**, and the rates shown are simply where the noise happened to peak.

Where it does appear it is unmistakable, and it is the downshifters:

FILE	FIRST THIRD	LAST THIRD	RATIO
Downshifter - Stutter Scream	4.1 Hz	25.0 Hz	6.08
Transition - Quick 05	2.6 Hz	7.2 Hz	2.75
Boom - Spacious	22.6 Hz	22.5 Hz	0.99
Downshifter - Groin Kick	4.8 Hz	2.4 Hz	0.49

2 to 25 Hz, both accelerating and slowing down, with the extremes at 0.22× and 6.08×. That is exactly a start speed and an end speed with an exponential glide between them.

So *Flutter* is **off by default**, and every preset that has it is one the measurement supports. *Start Speed* and *End Speed* are the pair; *Shape* chooses the modulator; *Target* decides whether it moves the level, the Air layer’s cutoff, or both — and on the cutoff is what most of the fluttering references are actually doing. *Smooth* rounds the modulator off, with a corner that follows the rate, so it means the same thing at 2 Hz and at 40.

Nothing above 25 Hz is thrown away

The rest of the Verdalis suite highpasses at 60 Hz, because in a field recording everything below that is traffic, ventilation and handling noise on the microphone rather than the thing being recorded.

Here it is the instrument. The *Highpass* default is 25 Hz and it is there to keep out DC and inaudible cone travel, nothing else.

The EQ

A three-band EQ on the summed output: a low shelf, a mid bell and a high shelf, with a highpass and a lowpass either side of them.

Low Gain is the control most of these sounds are actually mixed with. A cinematic impact is a bass decision, and it is a decision that changes with the track it is going into, which is why it is on the front panel rather than buried in the layers.

Mid Gain is how a transition is made to sit under a mix instead of in front of it. Cutting a few decibels around 700–900 Hz is the standard move and it is what *Brass Wall* does.

Variation

Every trigger draws its own values around the settings on screen: span, peak, cutoff, sweep, pitch, level, pan, flutter rate, hit tone and the struck layers' decays. The draws are gaussian and clipped at two sigma, so most triggers are near the setting and a few are noticeably not — which is what a folder of hand-made variants of one sound actually looks like.

At 0 the plugin is deterministic and repeats exactly, on any machine, at any sample rate and block size. The self-test checks that, and *Random Seed* pins the whole sequence when it is non-zero.

The plugin window

WhooshPact draws its own window with X11 and Cairo on Linux, and Cairo on the win32 backend on Windows. There is no toolkit dependency, which is what keeps the plugin one self-contained file. It is embedded in the host's own window through CLAP's GUI extension and repainted from the host's timer, and it resizes by zooming one layout, from half size to four times.

The window is the suite's, in WhooshPact's own impact-magenta theme. Every gesture fires a streak that crosses the header and lands in a flash at the far side. The streaks are deterministic in the gesture number, so a pinned *Random Seed* gives a repeatable picture as well as a repeatable sound.

Gestures

GESTURE	EFFECT
Drag a knob up or down	Change the value
Shift-drag	Fine control, a fifth of the travel
Double-click a knob, or right-click anything	Back to the default
Scroll wheel over a control	Step the value
Click a knob's value	Type a value in, with units
Click a selector's name	Open its list and pick a value
Click a selector's ◀ or ▶	Step one choice
Click the preset name	Open the preset browser
◀ or ▶ beside the preset name	Previous or next preset
SAVE	Save the current settings as a user preset
MIXER	Open the layer mixer

The line at the bottom of the window shows what the control under the pointer does — the same one-line explanation that appears in the *What it does* column of the parameter reference.

The mixer

Every layer's level lives on the panel that layer belongs to, which is right for editing one layer and wrong for balancing them against each other. **MIXER** opens the same controls arranged the other way round: one strip per layer, its level as a fader, its width as a slim slider under it. They are the plugin's own parameters, not a second set — moving a fader moves the knob on the panel, and the host sees the same automation it always would.

Four strips plus the output: **AIR**, **TONE**, **SUB**, **HIT**. Only Air and Tone carry a width, and none of them carries a pan: where a gesture sits is *Pan Start* and *Pan End* on the Motion panel, and that is a property of the whole gesture. A whoosh whose noise crossed the field while its sub stayed put would not read as one sound moving.

M mutes a layer and **S** solos one, which is how a preset gets built from nothing: solo one layer, get it right, bring the next one back. Several layers can be soloed at once.

Mute and solo are **not** parameters and are not saved with a preset. They work by holding a layer's level at the bottom of its range and remembering what it was, so while either is active the preset bar shows a **SOLO ON** or **MUTE ON** chip — click it to release every hold at once. **SAVE** releases them first, so a muted layer can never be written into a preset as a silent one, and so does loading a preset.

Parameter reference

Every range, default, unit and explanation below is read from the plugin's own parameter table when this manual is built, so it cannot drift from what the window and the host show.

There are 64 parameters, grouped the way the plugin window groups them. Every one is automatable and modulatable from the host, and every one can be typed in directly by clicking its value in the window. The name in `code` is the key used in preset files.

GESTURE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Type <code>type</code>	Accent, Boom, Braam, Downshifter, Impact, Transition	Transition	Which family's measured spectral profile the sound is given: the slope of its octave-band curve above 125 Hz, and how far its bottom two octaves stand over that curve. Both come from the 211-file reference library and both are relative to Transition, the broadest curve in it, which is therefore the neutral setting. It does not set the shape in time. That is Span, Peak, Hold, Rise and Fall, and those five reproduce every measured envelope contour in the library to within the library's own spread – so a type that moved them as well would only fight the preset that had just set them.
Blend <code>blend</code>	0.0 % ... 100.0 %	0.0 %	Crossfades from the chosen profile towards the next one, so the six measured shapes are a continuum rather than six settings. At 0 the sound is exactly the measured profile.
Span <code>span</code>	0.10 s ... 20.0 s	3.70 s	How long the whole gesture lasts. The plugin cannot know how long a note will be held, so this and not the note decides the shape: the rise, the middle and the fall are all fractions of it. Measured medians run from 3.6 s for an accent to 5.6 s for a braam, with the whole library between 1.7 and 12.1 s.
Peak <code>peak</code>	0.0 % ... 100.0 %	33.0 %	Where in the span the sound is loudest, as a fraction of it. This one number is what separates the families more sharply than anything else measured: a boom peaks at 0.024 and a downshifter at 0.024, an impact at 0.060, an accent at 0.043 – all of them at the very front – while a transition peaks at 0.328 and a braam at 0.148. So a hit is this near zero and a whoosh is this near a third.
Hold <code>hold</code>	0.0 % ... 100.0 %	10.0 %	How long the sound stays at full level after the peak, as a fraction of the span. Near zero for anything struck; a braam measures a plateau roughly a third of its span long, and a downshifter over half.
Rise <code>rise</code>	0.25 ... 8.00	1.80	The curvature of the approach to the peak. 1 is a straight ramp, above 1 holds back and then arrives, below 1 jumps and then creeps. Fitted to the measured transition contour, which rises as $t^{1.8}$.
Fall <code>fall</code>	0.25 ... 8.00	2.50	The curvature of the fall after the peak, on the same scale. Fitted to the measured impact contour, which is down to two thirds a tenth of the way through its fall and to a quarter by a third of it – about $t^{3.5}$.
Variation <code>variation</code>	0.0 % ... 100.0 %	35.0 %	How differently each trigger comes out. This is the reason the plugin exists: a sample plays the same transition every time and a track full of them reads as one sound repeated. Here every note draws its own span, peak, cutoff, pitch, pan, flutter rate and level around the settings on screen, and no two are alike. At 0 the plugin is deterministic and repeats exactly, which is what to use when a transition has to line up with a cut.

AIR

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Noise <code>noise</code>	White, Pink, Brown, Blue, Violet, Green	White	The colour of the noise the whoosh is made of. White is flat, Pink falls 3 dB/octave, Brown 6, Blue and Violet rise by the same, and Green is a broad mid-band emphasis around 500 Hz. Each is white noise with one filter on it, which is what those names have always meant.
Air Level <code>air_level</code>	-inf dB ... 6.00 dB	-6.00 dB	Level of the noise layer. On its own with everything else down this is the plain white-noise transition, which is the simplest thing the plugin makes.
Cutoff <code>air_cutoff</code>	40.0 Hz ... 18.00 kHz	2.40 kHz	Where the noise is filtered at the start of the gesture. The sweep moves it from here.
Sweep	-4.00	-1.60 oct	How far the cutoff travels across the gesture, in octaves. Negative closes down, positive opens up. The

air_sweep	oct ... 4.00 oct		default is fitted rather than copied. The library's transitions fall 0.80 octaves in spectral <i>centroid</i> from start to end – accents 0.69, impacts 0.63, downshifters 1.07, and braams alone rise, by 0.60. A centroid does not move as far as the cutoff that carries it, because what is left below the corner does not move at all: rendering this at -1.6 octaves measures -0.56 of centroid, which is where it sits.
Resonance air_reso	0.0 % ... 100.0 %	25.0 %	How much the filter rings at its corner – and, with it, how narrow the filter is. At 0 this is a plain lowpass and the whoosh is a wall of air closing down; at 1 it is a resonant bandpass and the whoosh has a pitch to follow. The two go together because that is how the ear hears it: a narrow moving band is a pitch, a broad one is wind.
Curve air_curve	0.25 ... 4.00	2.00	When in the gesture the sweep happens. 1 spreads it evenly; above 1 holds the cutoff still and then moves it late, which is what the measurement asks for – a transition's centroid is flat for its first third and then falls 0.8 octaves.
Tilt air_tilt	-100.0 % ... 100.0 %	0.0 %	Tilts the noise brighter or darker, +/-6 dB/octave about 1 kHz, on top of the colour chosen above.
Width air_width	0.0 % ... 100.0 %	60.0 %	How decorrelated the two channels of the noise are. At 1 they are independent, which is as wide as a noise source gets; at 0 they are the same signal and the whoosh is a point in the middle. The default is measured: the library's transitions have a median L/R correlation of 0.60, and this renders at 0.61.

TONE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Wave wave	Sine, Triangle, Saw, Square, Supersaw	Saw	The waveform of the pitched layer. Saw and Square are band-limited, because a braam glides upwards by an octave or more and a naive edge would alias audibly on the way. Supersaw is the saw stack detuned against itself.
Tone Level tone_level	-inf dB ... 6.00 dB	-inf dB	Level of the pitched layer. Off in the default patch – a plain noise transition has no pitch in it – and the loudest thing in every braam.
Pitch tone_pitch	20.0 Hz ... 2.00 kHz	49.0 Hz	Where the pitched layer starts. The default is measured: the strongest partial below 200 Hz in the loudest half second of a braam is 47 Hz, of a boom 49, of a downshifter 41 – all of them within a tone of G1. The note played transposes this.
Glide tone_glide	-48.00 st ... 48.00 st	-12.00 st	How far the pitch travels across the gesture, in semitones. Down is a downshifter, up is a braam opening out – the library's braams rise a measured 0.6 octaves in centroid and their low partial by more than an octave.
Detune tone_detune	0.0 % ... 100.0 %	25.0 %	How far the three oscillators of the stack are pulled apart. What turns one note into a section.
Width tone_width	0.0 % ... 100.0 %	60.0 %	How far the detuned oscillators are spread across the stereo field.

MOTION

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Pan Start pan_start	-100.0 % ... 100.0 %	-60.0 %	Where the gesture sits at its beginning. With Pan End this is the passby: a whoosh that crosses the listener rather than sitting in front of them. The library's transitions measure an L/R correlation of 0.60 against a boom's 0.89, which is most of what makes one feel like movement.
Pan End pan_end	-100.0 % ... 100.0 %	60.0 %	Where it has arrived by the end. Equal to Pan Start for a sound that does not travel.
Space space_amount	0.0 % ... 100.0 %	22.0 %	How much reverberant field there is around the gesture. Cinematic impacts live or die by this: the tail is most of what makes one sound big.
Size space_size	0.0 % ... 100.0	45.0 %	How big that space is.

	%		
Damping space_damping	0.0 % ... 100.0 %	50.0 %	How quickly the space eats the top end.
Reverb Width space_width	0.0 % ... 100.0 %	80.0 %	How wide the reverberant field is. Narrow keeps the tail behind the sound; wide puts it all around.

SUB

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Sub Level sub_level	-inf dB ... 6.00 dB	-inf dB	Level of the low layer, off by default so that a patch opts into it. A boom is almost nothing else: the library measures 97 per cent of a boom's energy and 98 per cent of a downshifter's below 100 Hz, against 62 per cent for a transition – but a transition that has to sit under a mix often wants no low end at all.
Sub Pitch sub_pitch	15.0 Hz ... 200 Hz	49.0 Hz	Where the low layer starts. 49 Hz is the measured median fundamental of the booms, which is G1.
Drop sub_drop	-48.00 st ... 12.00 st	-12.00 st	How far the low layer falls across its decay, in semitones. A boom that does not drop sounds like a note; one that drops sounds like something hitting the ground.
Decay sub_decay	20.0 ms ... 8.00 s	700 ms	How long the low layer takes to fall 20 dB. Measured 0.63 s at 63 Hz for the booms, 0.78 for the impacts and 1.9 for the downshifters, against 0.35 for the accents – the low decay is most of what tells those apart.
Drive sub_drive	0.0 % ... 100.0 %	20.0 %	Saturation on the low layer alone. A pure sine at 45 Hz disappears on small speakers; the harmonics this adds are what carry it there.
Click sub_click	0.0 % ... 100.0 %	30.0 %	A very short upward pitch snap at the start of the low layer. It is what makes a boom read as an impact rather than as a note fading in.

HIT

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Hit Level hit_level	-inf dB ... 6.00 dB	-inf dB	Level of the transient. Off in the default patch, and the whole point of every impact and accent preset.
Tone hit_tone	60.0 Hz ... 12.00 kHz	900 Hz	Where the transient's energy sits.
Decay hit_decay	5.00 ms ... 3.00 s	220 ms	How long the transient takes to fall 20 dB. The library's accents measure 0.32-0.50 s across the octaves above 500 Hz and its impacts 0.27-0.53, which is a short sound with a long room on it rather than a long sound.
Noise hit_noise	0.0 % ... 100.0 %	60.0 %	Morphs the transient from a cluster of three inharmonic resonators towards plain filtered noise. Low is metal, high is debris.
Body hit_body	0.0 % ... 100.0 %	50.0 %	How much low end the transient carries under its tone. Without it a hit is all crack and no weight.
Time hit_time	0.0 % ... 100.0 %	33.0 %	Where in the span the transient lands, as a fraction of it. At the peak for an impact; at the end of a whoosh for a whoosh-hit, which is what the library files under Transition - Whoosh Hit.

FLUTTER

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Flutter flutter_depth	0.0 % ... 100.0 %	0.0 %	How deeply the gesture is chopped up. Zero is off, and off is the default – the library is mostly not fluttering. Measuring it says so plainly: the fraction of the in-band modulation energy sitting at a single rate is 0.02 to 0.05 in five of the six families, which is noise with no rate at all. Where it does appear it is unmistakable, and it is the downshifters that have it: Downshifter - Stutter Scream measures 0.12 and accelerates from 4.1 to 25.0 Hz across its span. So this is a deliberate effect with a measured range rather than something every preset gets.
Start Speed flutter_start	0.50 Hz ... 60.0 Hz	6.00 Hz	How fast the chopping is at the beginning of the gesture. Measured 2.6 to 25.0 Hz across the references that have any.
End Speed flutter_end	0.50 Hz ... 60.0 Hz	14.0 Hz	How fast it is by the end. The rate glides exponentially between the two, which is what a single pair of numbers can do that an LFO cannot. The measured end-to-start ratios run from 0.22 to 6.08 – both accelerating and slowing down happen, and the downshifters' median is a doubling.
Shape flutter_shape	Sine, Triangle, Square, Ramp, Random	Sine	What the chopping looks like. Sine is a tremolo, Square a gate, Ramp a repeated fall, Random a new level each cycle.
Target flutter_target	Level, Filter, Both	Level	Whether the chopping moves the level, the Air layer's cutoff, or both. On the cutoff it is the filter fluttering rather than the amplitude, which is what most of the downshifters in the library are doing.
Smooth flutter_smooth	0.0 % ... 100.0 %	35.0 %	Rounds the modulator off. A hard square at 20 Hz clicks on every edge; this is how much of that is taken away.

EQ

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Highpass highpass	15.0 Hz ... 1.00 kHz	25.0 Hz	Permanent highpass, and set low on purpose. The rest of the suite starts at 60 Hz because below that a field recording is traffic and ventilation; here the sub-60 Hz band is the instrument. 25 Hz keeps the DC and the inaudible cone travel out and nothing else.
Lowpass lowpass	500 Hz ... 22.00 kHz	18.00 kHz	Permanent lowpass, for taking the fizz off the top of a noise layer.
Low Freq eq_low_freq	30.0 Hz ... 400 Hz	90.0 Hz	Corner of the low shelf.
Low Gain eq_low_gain	-18.00 dB ... 18.00 dB	0.00 dB	Low shelf gain. This is the one control most of these sounds are actually mixed with: a cinematic impact is a bass decision.
Mid Freq eq_mid_freq	150 Hz ... 6.00 kHz	900 Hz	Centre of the mid bell.
Mid Gain eq_mid_gain	-18.00 dB ... 18.00 dB	0.00 dB	Mid bell gain. Cutting here is how a transition is made to sit under a mix instead of in front of it.
High Freq eq_high_freq	1.00 kHz ... 16.00 kHz	5.00 kHz	Corner of the high shelf.
High Gain eq_high_gain	-18.00 dB ... 18.00 dB	0.00 dB	High shelf gain.

ENVELOPE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Attack attack	0.20 ms ... 2.00 s	1.00 ms	Fade-in applied on top of the gesture. Short by default: the gesture's own Rise is what shapes the front of the sound, and this is only here to take the edge off it.
Decay decay	5.00 ms ... 20.00 s	2.00 s	Fall from the initial level to the sustain.
Sustain sustain	0.0 % ... 100.0 %	100.0 %	Level held while the note is down. At 1 the gesture plays out exactly as its own shape says, whatever the note does.
Release release	5.00 ms ... 20.00 s	300 ms	Fade-out after the note is let go. This is how a gesture is cut short: the span decides how long the sound wants to be, and letting go early with a short release ends it anyway.
Velocity To Level vel_to_level	0.0 % ... 100.0 %	50.0 %	How much velocity sets the level.
Velocity To Tone vel_to_tone	0.0 % ... 100.0 %	30.0 %	How much velocity opens the Air layer's filter and hardens the transient – which is what hitting something harder actually does.

OUTPUT

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Output Gain gain	-inf dB ... 12.00 dB	-4.00 dB	Output level.
Drive drive	0.0 % ... 100.0 %	12.0 %	Saturation on the summed output. The library's crest factors run from 7.7 dB for the downshifters to 19.0 for the accents, and the low end of that is a sound that has been driven hard.
Max Voices max_voices	1 ... 16	8	How many gestures may overlap. One is monophonic, which is what a transition usually wants; more lets a new hit land while the last one is still ringing.
Random Seed seed	0 ... 999	0	Non-zero gives the same sequence of gestures every time, which is what makes a rendered transition reproducible. Zero is always different.

The preset library

Thirty-one presets: five in each of the six families, plus one deliberately thin accent. Generated from the preset files themselves.

31 factory presets ship with the plugin. They are installed beside the binary and exposed through CLAP preset discovery, so they appear in the host's own browser as well as in the plugin window's.

Analog Fall

Two octaves of falling saw with the flutter gliding up underneath it, from 3 Hz at the top to 14 by the bottom. The measured median of the family is a doubling; this is closer to five times, which is what the ear reads as falling faster than it is. (*downshifter, analog, fall*)

Annihilation

The same horn pulled apart: detune at the top of its range and the stack spread fully across the field, so it reads as a section rather than a note. (*braam, wide, dense*)

Bad Guy

Square rather than saw, and low: the hollow odd-harmonic version of a braam, slow to arrive and slow to leave. The square at 37 is already the bottom of the sound, so there is no sub under it – adding one moved the low fraction by 0.02 and nothing else. (*braam, square, menace*)

Big Horn

The archetype: a detuned saw stack that swells for a second and a half and opens as it goes, which is the one thing the measurement says a braam does and nothing else in the library does. (*braam, horn, swell*)

Blast Off

An impact that leaves something behind: the hit lands first and the noise layer opens *upwards* after it, which is the one place in this family the measured downward sweep is deliberately inverted. (*impact, riser, tail*)

Bottomless

The deepest thing in the library: a sub at 42 Hz falling most of an octave across four seconds, with a large soft room under it and nothing above 1 kHz. (*boom, deep, long*)

Brass Wall

The loudest of the five: saw stack, sub and noise all up, driven into the output stage so the crest factor comes down to where the references sit. (*braam, loud, brass*)

Deep Impact

The cinematic one: a hard transient on top of a long low decay, with the 0.78 s at 63 Hz and 0.21 s at 8 kHz the library measures built in as the difference between Sub Decay and Hit Decay. (*impact, cinematic, deep*)

Deep Sonar

A boom with a note in it: the tone layer at the same pitch as the sub and barely detuned, so the decay has a pitch to follow all the way down. (*boom, tonal, sonar*)

Deep Whoosh

The same gesture an octave and a half down: brown noise, a low corner and a little sub under it. The brown noise and the 900 Hz corner do most of the work – the sub is trimmed to sit behind them rather than replace them. (*transition, deep, low*)

Digging Deep

No pitched layer at all: the sub and the noise together, with the flutter on the cutoff rather than the level. That is what most of the references are actually doing – the filter is fluttering, not the amplitude. (*downshifter, filter, deep*)

Dry Snap

The quiet end of the library, and the accent counterpart to Simple Whoosh: one short noise stab and nothing at all under it – no sub, no transient, barely any room. Deliberately outside the family's 0.69 below 100 Hz, for a cut where the mix has no space left for a low layer. (*accent, thin, dry*)

Gate Closed

A heavy door: a dull thud with one metallic mode left ringing on top of it, and the room doing most of the work afterwards. No sub layer – the hit's own body carries the weight, and a sub under it measured 0.90 below 100 Hz against the family's 0.69. (*accent, thud, metal*)

Glass Break

All top and no bottom – no sub layer at all, and a 180 Hz corner under it: a bright noise burst with the hit layer pushed fully into debris, and violet noise over it so the shards keep coming. (*accent, bright, debris*)

Grind To A Halt

The other direction, which the measurement says happens just as often: the flutter starts fast and slows to almost nothing, so the gesture reads as something losing power rather than gaining it. (*downshifter, decelerate, grind*)

Head-on Collision

Loud and broken: everything driven, the transient pushed fully into debris, and the widest stereo picture of the five. (*impact, debris, loud*)

Jumpscare

The one accent with a run-up: a very short rising whoosh into a hard hit, so the ear has just enough time to know something is coming. (*accent, riser, scare*)

Low Thud

An impact with the top taken off: nothing above 1.2 kHz, a short body and a dry room. What an impact sounds like from the other side of a wall. (*impact, dark, dry*)

Metallic Collision

Two large metal things meeting: the hit layer's inharmonic cluster with the noise mix low, so the modes ring past the transient that excited them. (*impact, metal, ring*)

Passby

Short, fast and fully across the field: a narrow resonant band sweeping hard, which is what makes the ear hear something go past rather than fade. No low layer – the 90 Hz corner would take it out again anyway. (*transition, passby, movement*)

Retro Drop

A square wave, a fast hard gate and no room to speak of: the one preset in the family that sounds like a machine rather than like a collapse. (*downshifter, square, retro*)

Reveal

A long slow opening: the whole span is the rise, the filter opens instead of closing, and the gesture stops the moment it arrives. The one transition shape a sample library files separately and the measurement bears out. (*transition, riser, reveal*)

Rising Dread

A braam used as a riser: a full octave of glide, the noise layer opening with it, and no plateau at all – it arrives at the top and stops. (*braam, riser, glide*)

Shattered Earth

A boom that breaks something: the full sub with the hit layer pushed into debris over it, landing a moment after the impact rather than with it. (*boom, debris, impact*)

Simple Whoosh

The default patch, and the one everything else is built out of: white noise through one swept filter, crossing from left to right, and nothing else at all – no sub, no tone, no transient. Every measured quantity in the gesture is the library's median for the family. (*transition, whoosh, default*)

Steel Stab

A struck metal accent: the hit layer's three inharmonic modes with almost no noise over them, a short bright body and a room that lets them ring out. (*accent, metal, stab*)

Stutter Scream

Fitted directly to **Downshifter - Stutter Scream**, which accelerates from 4.1 Hz to 25.0 Hz across its span and is the most strongly periodic thing in the whole 211-file library. A square gate on the level, so it chops rather than wobbles. *(downshifter, stutter, gate)*

Underground Roll

Driven rather than deep: the sub saturated hard so it carries on a small speaker, with a slow flutter under it that keeps the decay moving. *(boom, driven, roll)*

War Drum

The low end of an accent rather than its edge: a struck drum with the click turned well up, tuned an octave under the stabs and given nothing above 3 kHz. *(accent, drum, low)*

Whisper

A boom a long way off: quiet, almost no click, and mostly the room it happened in. Useful under dialogue, where a loud one is unusable. *(boom, distant, quiet)*

Whoosh Hit

A whoosh that lands: the noise layer rising to a peak two thirds of the way in and the transient timed to arrive exactly there. The library files fifteen of these under Whoosh Hit, and the timing is the whole trick. *(transition, whoosh, impact)*

Using it in a host

Notes

One note is one gesture. Velocity does two things: *Velocity To Level* sets how loud it is, and *Velocity To Tone* opens the Air layer's filter and hardens the transient, which is what hitting something harder actually does.

The played note transposes the Tone and Sub layers from middle C. A preset plays as it was fitted at middle C, an octave down at C2, and so on. The Air layer does not track the note — noise has no pitch to transpose.

Max Voices is how many gestures may overlap. One is monophonic, which is often what a transition wants; more lets a new hit land while the last one is still ringing. Turning it down while gestures are sounding releases them rather than cutting them off.

Automation and modulation

Every parameter is automatable and modulatable. Note that automating a parameter *during* a gesture works, but several of the gesture's own numbers — *Span*, *Peak*, the variation draws — are read once when the note starts, so changing them mid-gesture affects the next trigger rather than the current one.

Reproducible renders

Set *Variation* to 0 for an exactly repeatable gesture, or leave *Variation* where it is and set *Random Seed* to a non-zero number for a repeatable *sequence* of different ones. Either way the output is bit-identical across runs, machines, sample rates and block sizes, which is what makes a bounced transition match the one that was auditioned.

CPU cost

Bounded by *Max Voices*. Each voice carries two noise generators, a state-variable filter per channel, a tilt pair, three oscillators, a sine and a resonator cluster; everything that costs a transcendental is recomputed once per 32 samples and interpolated. Sixteen simultaneous gestures is a small fraction of one core.

Rendering and tails

The plugin reports a tail to the host that covers the longest a gesture can run: its span plus the widest a varied span can get, plus the struck layers' decays and the reverb. A bounce that stops at the last note-off will still have the whole tail in it if the host honours `CLAP_EXT_TAIL`, which most do.

Headroom

The factory presets are trimmed to about -8 dBFS peak at *Random Seed* 7, with *Variation* at its preset value. The output stage is the suite's soft clipper, so driving it harder saturates rather than clips — but a gesture is a peaky thing and the library's crest factors run from 7.7 dB to 19.0, so leave room.

What is still being fitted

The ear has not had its turn. The suite's rule is ear first, then measurement, and this plugin currently has the second without the first. The measurements say the shapes and the balances are right; they cannot say whether a braam sounds like a braam. A/B rendering against the references is the next piece of work.

The Hit layer is enveloped, not ringing. Its three modes are excited by enveloped noise and their decay is that envelope rather than their own Q, which is why a long *Hit Decay* reads as noise held open instead of as metal ringing down. A proper modal resonator with a decay rate per mode would fix it, and would belong in the suite's shared DSP rather than here.

There is no debris. The measured impacts are often not one event at all — *Impact - Bits & Pieces* and *Impact - Complex Wreck* are scatters of small events over a second or more. The plugin has no way to make one.

No tempo sync. *Span* is in seconds and the flutter rate is in hertz. A transition is almost always cut to the bar.

Crest factor. The downshifter references measure 7.7 dB; the presets render at 11–13. Those are limited masters and the output stage here is a soft clipper, so the gap is expected rather than surprising — but it is a gap.

Reference and limits

Where the model comes from

211 production sounds in six families — 35 accents, 16 booms, 17 braams, 13 downshifters, 24 impacts and 106 transitions — measured with the scripts in `tools/analysis/`, which report:

- **the shape in time:** span, peak position, the 10–90 % rise and fall, and the median envelope contour of each family on a normalised time axis
- **the sweep:** the spectral centroid at the start, the middle and the end, and its contour
- **the spectrum:** octave-band levels from 31.5 Hz to 16 kHz, the fraction below 100 Hz, and the per-band decay to –20 dB
- **the pitch:** the strongest partial below 200 Hz in each gesture's loudest half second
- **the flutter:** the rate, depth and periodicity of the amplitude modulation in the first and last third of each gesture

The reference material is not in this repository and is not ours to redistribute.

One measurement trap is recorded in `tools/analysis/README.md` because it cost an hour: a 4 ms RMS envelope has its Nyquist at 125 Hz and will happily track the *carrier* of a sound that has descended to 30 Hz, which is exactly what a downshifter does. The first run reported flutter at 34 Hz for files with none.

What it does not do

- **It is not a sampler and will not reproduce a specific reference.** It is fitted to what those sounds *are* statistically, not to any one of them.
- **No tempo sync,** no host transport, no reverse mode.
- **The Air layer does not follow the note.** Noise has no pitch, so only Tone and Sub transpose.
- **Blend wraps from the last family to the first.** Blending past Transition arrives at Accent, which is alphabetical rather than musical.
- **A gesture cannot be retriggered from where it was.** A new note is a new gesture; there is no legato.