



VERDALIS

Natural sound synthesis

NightLife

A synthesised night



Version 0.3.0

CLAP instrument for Linux and Windows

Part of the Verdalis Plugin Suite

2026-10-04

Contents

Introduction

- What is in this manual

Installing

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

A first sound

- Four things to try first

How it works

- The contour is the animal
- The term count is per second, not per call
- The four contour controls
- The valve
- The partials: reaching for the recording
- The tract
- A phrase, and a pack
- The chorus is a pulse train
- Why the pond keeps time
- The insects
- The bed
- Distance

The plugin window

The parameter reference

- Call
- Voice
- Phrase
- Pack
- Chorus
- Insects
- Bed
- Place
- Filter
- Envelope
- Output

The preset library

Introduction

NightLife is a CLAP instrument that generates a night. A wolf howling across a valley, an owl in a wood, a fox screaming in a field, a loon over a lake, a pond full of frogs, a field of crickets, and the quiet that all of it sits on. Every sound is computed while the plugin plays — nothing is played back, and no two takes are alike.

A night is not one texture, so NightLife is not one mechanism. It has four layers and they are built from three different kinds of measurement, because the things they model are different kinds of thing.

A call is its frequency contour. Not a pitch envelope on an oscillator — the actual shape the pitch traces. The plugin does not compute those shapes, it *measures* them: 48 contours were extracted from real recordings of wolves, owls, foxes and a loon, fitted as formulas, and they drive the oscillator directly. *Contour* chooses which one. A wolf's vibrato is not a control here; six of the eight wolf contours carry one at 2 to 6 Hz because the wolves that were recorded did.

A croak is not a contour. A frog is a pulse train through a body resonance — 10 to 69 pulses a second, and the rate is the species. Eight recordings were measured for their rate, pulse count, both resonances and envelope, and *Croak* walks across them.

A chorus of frogs is more regular than random. Every other event in the Verdalis suite is spawned as a Poisson process. Measured, a frog chorus is not one: its arrivals are *under*-dispersed, strongly so over a few seconds. So the pond here keeps one rhythm and each frog takes a place in it, which is what a chorus of real frogs does.

The insects and the night itself are noise, measured. A cricket band is narrow, high and steady; the night behind everything is the third-octave curve of the quietest parts of 45 field recordings, solved onto a filterbank.

NightLife 0.3.0 is the first release. Nothing before it was published, and no preset or saved state from an earlier build exists.

What is in this manual

Installing is the boring part. *A first sound* is five minutes. *How it works* is what the four layers actually do and why they are built the way they are. *The parameter reference* lists every control with its range and default, and is generated from the plugin's own table, so it cannot drift from the build. *The preset library* is the same for the factory presets.

Installing

NightLife is one self-contained `.clap` file plus its `presets` folder, and a `.vst3` bundle beside it. There is nothing to register and no runtime to install.

Linux

Copy the plugin folder into your CLAP directory:

```
~/clap/NightLife/NightLife.clap
~/clap/NightLife/presets/
```

and, for the VST3, `~/vst3/NightLife.vst3`. A system-wide install goes in `/usr/lib/clap` and `/usr/lib/vst3`.

Windows

```
C:\Program Files\Common Files\CLAP\NightLife\
C:\Program Files\Common Files\VST3\NightLife.vst3
```

Copy the whole `NightLife` folder so that `NightLife.clap` and `presets\` sit inside it, and rescan.

Hosts

Any CLAP host: Bitwig Studio, Reaper, Qtractor, or the reference `clap-host`. NightLife appears as an instrument under the vendor **Ravetracer**. In a VST3 host it appears the same way; the plugin's own preset browser lists the factory presets there too, though the host's own browser will not.

Your own presets

SAVE in the plugin window writes to your user preset folder:

```
| Linux | $XDG_CONFIG_HOME/NightLife/presets, or ~/.config/NightLife/presets | | Windows |
%APPDATA%\NightLife\presets |
```

They are plain text — parameter names and values — and can be edited, copied between machines and kept in version control.

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 `.nightlifepack` — 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

```
./install.sh          # configure, build, self-test, install to ~/.clap
./install.sh --vst3    # and the VST3
```

A first sound

1. Load NightLife on an instrument track.
2. Choose **Wolf Valley** in the preset bar.
3. Hold a long note. Middle C is the pitch the preset was written at; every other key transposes the whole night.

The note does two things. It fires one phrase immediately — that is *Shot Level* — and it brings up everything that happens unprompted while it is held: the pack, the chorus, the insects and the bed. Let go and the layers fade over *Release*, but a call already under way always finishes. A howl is never cut in half.

Four things to try first

Turn *Distance* up. It is the most effective control in the plugin. A wolf at the far end of a valley is most of what the word “howl” means, and distance is what makes one — the level falls, and the air takes the top off the call.

Open the MIXER. Four layers is one more than anything else in the suite has, and balancing them is most of what building a night from scratch is. Each fader has a mute and a solo; solo the chorus, get it right, bring the rest back.

Turn *Answer up on Pack Answering*. One animal calls and another replies from somewhere else a moment later. That behaviour, and not the voice, is what makes a group of wolves sound like a group.

Turn *Regularity* down on Pond Chorus. At zero the pond becomes a Poisson process — the way every other plugin in the suite spawns events — and it is immediately audible as *wrong*: the croaks clump. That difference is a measurement, not a preference.

How it works

The contour is the animal

A wolf howl is one pitch held for a second or more, with a slow vibrato on it and a fall at the end. Everything that makes it a wolf rather than a synthesiser is in how that line actually moves, and the line is not a curve anybody would guess: it wanders, it breaks, it is not symmetric, and its vibrato is not a sine.

So NightLife measures it. `tools/analysis/contours.py` pulls the instantaneous frequency and level of every well-isolated call in the reference library out of the recording, fits each as a cosine series in normalised call time, clusters them per caller and keeps the medoids — a medoid rather than an average, because averaging two contours that move out of phase gives a smooth glide, and a smooth glide is exactly what a synthesised howl sounds wrong as.

No audio is stored. A contour is a formula: the same thing as

```
f(x) = -260 sin(-pi x / 5200) + 1740
```

read off a spectrogram by hand, with one term per 18 ms of call instead of one term in total, because a real call turns direction far more often than once.

Eight archetypes per caller, 48 in all, 43 kB.

The term count is per second, not per call

This is the one place NightLife had to depart from ChirpParade, which does the same thing for birds. Its fit budget is a fixed 96 terms per syllable, which is right when everything is between 12 and 900 ms long. Here the calls run from a 75 ms pip to a 3.1 s whinny, and a fixed budget is then a hidden duration filter: at 96 terms the wolf reproduces 97 % of its measured motion and the scops owl 253 % of its — the series ringing between the points it was fitted at.

55 terms per second of call holds all six callers between 77 % and 115 %. Each archetype carries its own count.

The four contour controls

Contour chooses which measured call, walking across that caller's eight from the lowest-sitting to the highest. *Sweep* scales how far it travels — 100 % is the measured curve exactly. *Detail* is how much of its fine motion survives; at the top the measured curve passes through, and turning it down smooths the contour towards the glide a description of a howl would have given. *Skew* bends the call's own time axis.

Length scales the archetype's **own** measured duration rather than replacing it, so a caller's spread of call lengths survives. The wolf archetypes run 311 ms to 2.8 s; stretching them all to one target is what a summary statistic would do.

The valve

Air passes through a larynx only while the folds are apart. *Voice* is the fraction of each cycle they are shut. At the bottom the valve never closes and what comes out is very nearly a pure tone — which is what a hoot is, measured at one harmonic and -34 dB of roughness. Close it and the airflow becomes a one-sided pulse with the harmonic stack a scream has, evens as well as odds. That is not a filter, and there is nothing between the two settings because there is nothing between them in the animal: a symmetric oscillator has no even harmonics at all, whatever it is driven with.

Rasp is irregularity of that closure, cycle by cycle. The break at the top of a howl and a vixen's scream are both this rather than more harmonics — the fox group measures -23 dB of roughness against every other caller's -32 to -34, and no amount of extra harmonics is the same thing as noise.

The partials: reaching for the recording

Partials crossfades the synthetic valve into the archetype's **own measured balance** between its first six harmonics, across the call. It is the one control that reaches for the recording itself: a fox opens and closes its harmonic stack within a single scream, and a fixed valve through a fixed tract cannot do that at all.

It is scaled by how much of that call's energy the harmonic measurement actually accounted for, so an archetype with a second animal in it does not pretend to know.

The tract

The tube above the larynx, closed at one end and open at the mouth, resonating at $c/4L$. *Throat* is its length — 19 cm puts the resonance at 450 Hz, where the library's hoots and howls sit. *Muzzle* opens the mouth: it raises the resonance, damps it, and makes it follow the pitch, which is what an animal that opens its jaw as it rises actually does. A wolf does exactly that.

A phrase, and a pack

Calls is how many calls a phrase holds and *Call Rate* how fast they follow each other, onset to onset. The library's median is three calls about a second apart, and the spread is the whole point: a tawny owl's phrase is two notes, a wolf's is one held note, a barred owl's eight, a scops owl's fifty-seven of the same pip.

Legato stretches each call towards filling its slot. It defaults to 32 % because that is the measured share of consecutive calls in the library that touch. ChirpParade measures 70 % for birds and defaults accordingly; a night is not a dawn chorus.

The pack is individuals, not a rate. *Animals* draws that many, each with its own pitch, position, distance, voice, tempo and level, held for as long as the note is. *Pitch Spread* is wide by default because wolves in a chorus avoid each other's pitch — which is why a pack sounds like more animals than it holds. *Answer* is how often one replies to another shortly after, and it is most of what makes a pack sound like a conversation rather than a random process.

The chorus is a pulse train

A croak is not a pitch contour and is not modelled as one. The frog drives its vocal sac in bursts, and what arrives is a train of pulses, each ringing the sac and the mouth. *Pulse Rate* is the species — it is the first thing a field guide lists, and it is what a listener hears as the difference between a knock, a creak and a rattle. Measured across the library: 10 to 69 Hz.

Croak Pitch sets the first resonance. The second one moves with it, at the ratio that row of the table measured — about half the first, across the whole library. A croak with only one resonance is a beep.

Why the pond keeps time

The Verdalis suite spawns events as a Poisson process everywhere else. Measured, a frog chorus is not one. The statistic is the Fano factor — the variance of the arrival count in a window over its mean, which for a Poisson process is exactly 1.0 at every window:

WINDOW	50 MS	250 MS	1 S	4 S
a measured chorus	0.90	0.81	0.59	0.30
a Poisson process	1.00	1.00	1.00	1.00

A frog chorus is *more* regular than random. And giving each frog a clock of its own does not reproduce it: enough independent clocks add up to a Poisson process however regular each one is. The regularity belongs to the **pond**. So the chorus keeps one period and each frog takes a place in it, which is also what real frogs do — they call in turn.

Regularity is the control between the two. At the top the pond has a rhythm; at the bottom it is the Poisson process the rest of the suite uses, and the difference is audible immediately.

The insects

A cricket is a narrow band of noise with a trill on it. There is no pitch contour and no body resonance in one, and none is modelled: *Insect Pitch* is the carrier, *Insect Band* its width, *Trill Rate* and *Trill Depth* the pulsing, *Shimmer* how much the individuals differ.

This is the weakest measurement in the plugin and it is worth saying so plainly. The reference library has no recording of insects on their own — what is measured is the cricket band *behind* a scops owl and behind two frog choruses, isolated by being narrow, high and steady where everything else in those files is none of the three. Two recordings carry a prominent trill rate and they disagree by an octave, 33 and 49 Hz, which is plausibly a katydid and a cricket.

The bed

The night behind everything: wind in leaves, distant water, a chorus too far away to be individual animals. Its shape is measured — the third-octave curve of the quietest third of the frames of all 45 references — and solved onto an eight-band octave filterbank.

Bed Tilt turns that curve about 500 Hz, where it peaks. *Bed Motion* drifts its level slowly, so it breathes rather than sitting still.

Distance

Distance does two things at once and both are physical: the level falls with distance, and the air absorbs the top of the sound. *Depth* is how much the animals differ in it, which is most of what makes a night sound deep rather than wide. *Air* is how much of the high end survives — cold dry air keeps more.

The plugin window

NightLife draws its own window with X11 and Cairo on Linux, and Cairo on the Win32 backend on Windows. There is no toolkit.

Drag a knob to edit it, double-click to reset it to the default, shift-drag for fine control, and click a value to type one. An enum chip steps through its choices when clicked on its left or right half.

The preset bar holds the browser, *SAVE*, and *MIXER*. The mixer is an overlay with one strip per layer: its level as a fader, a mute and a solo. **Mute and solo are not parameters** — a preset is a set of parameter values, and a mute that reached the parameters would be saved as a layer that comes back silent. A *SOLO ON* or *MUTE ON* chip stays in the preset bar for as long as a hold is active, because a forced-down layer that looks like a saved one is the trap.

The header draws a night sky over a treeline, and every call sends a ring out from a point on the horizon. You hear a wolf long before you could see one, and what reaches you is a wavefront.

The parameter reference

There are 75 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.

CALL

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Pitch <code>pitch</code>	60.0 Hz ... 6.00 kHz	893 Hz	Where the call sits. The fundamental across the library runs 138 Hz to 3.1 kHz with a median of 893, which is the default: a wolf and a big owl at the bottom, a vixen's scream at the top. The contour is measured about its own centre, so this transposes it – and because the default is the library median, every caller at the default calls in its own register.
Contour <code>contour</code>	0.0 % ... 100.0 %	50.0 %	Which call. Each caller carries eight contours measured off real recordings – the medoids of its clustered calls – and this walks across them, lowest-sitting first. It is the most important control in the plugin: a call's identity is its frequency contour, and these are measured curves rather than a shape derived from summary statistics. The wolf's vibrato is in them, at the 2 to 6 Hz the recordings actually carry.
Detail <code>detail</code>	0.0 % ... 100.0 %	100.0 %	How much of the contour's fine motion survives. At 100 % the measured curve passes through; lower it and the contour smooths towards a glide – which is what a howl synthesised from a description rather than a measurement sounds like, and the reason this pipeline exists.
Sweep <code>sweep</code>	0.00 % ... 300.00 %	100.00 %	Scales how far the contour travels. 100 % is the measured curve exactly; below that it flattens towards a held note, above it exaggerates. It goes past 100 % on purpose: a wolf holding one note is a real howl and so is one that swings a fifth, and the archetypes span 0.04 to 1.4 octaves.
Length <code>length</code>	40.0 ms ... 6.00 s	490 ms	How long one call lasts. The archetypes run 75 ms to 3.1 s with a median of 490, and each plays at its own measured duration scaled by this – so a caller's spread of call lengths survives instead of every curve being stretched to one target.
Skew <code>skew</code>	0.0 % ... 100.0 %	50.0 %	Bends the call's own time axis: below a half it crowds the contour towards the start, above it towards the end. 0.5 plays the measured curve at its measured pace, and the asymmetry the library shows – a 137 ms rise against a 278 ms fall – is already in it.
Jitter <code>jitter</code>	0.0 % ... 100.0 %	10.0 %	How much the pitch wanders off its own contour. A larynx is not a synthesiser: without this every call of a phrase is identical, which no recording in the library is.
Vibrato <code>vibrato</code>	0.0 % ... 100.0 %	0.0 %	Extra vibrato, on top of whatever the measured contour already does. It defaults to nothing for a reason: six of the eight wolf archetypes carry a measured vibrato of their own, between 2.1 and 6.4 Hz, so this is for pushing further rather than for supplying what is missing.
Vibrato Rate <code>vibrato_rate</code>	0.50 Hz ... 20.0 Hz	3.40 Hz	Rate of that extra vibrato. The default is the median measured across every archetype that has one, 3.4 Hz.

VOICE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Caller <code>caller</code>	Wolf, Owl, Screech, Scops, Fox, Loon	Wolf	Which animal. Each entry chooses that caller's own set of measured contours and biases pitch, length, harmonic richness, roughness and call rate together, because in the references they move together. Every figure is the median of the recordings of that animal; nothing here is invented.
Voice <code>voice</code>	0.0 % ... 100.0 %	30.0 %	How much of each cycle the larynx is shut. Air passes only while the folds are apart, so this is not a filter: at the bottom the valve never closes and a pure tone comes out, which is what a hoot nearly is, and closing it makes the airflow a one-sided pulse with the harmonic stack a scream has – evens as well as odds, which a

			symmetric oscillator cannot produce at all.
Breath breath	0.0 % ... 100.0 %	6.0 %	Turbulent air past the folds. It is also what starts the oscillation: the noise is injected into the oscillator rather than added to its output, which is why a call's onset is never twice the same. The chosen caller scales it, from the roughness its own recordings measure – a fox is 10 dB rougher than everything else in the library.
Throat throat	1.00 cm ... 60.00 cm	19.00 cm	Length of the tube above the larynx, which resonates at $c/4L$. 19 cm puts the resonance at 450 Hz, which is where the library's hoots and howls sit; a fox is a much shorter tube and a wolf a longer one.
Muzzle muzzle	0.0 % ... 100.0 %	50.0 %	How far the mouth is open. It raises the tube's resonance, damps it, and makes it follow the pitch – which is what an animal that opens its jaw as it rises actually does, and a wolf does exactly that.
Formant formant	0.0 % ... 100.0 %	55.0 %	How much of that resonance reaches the output. At zero the voice is the bare valve; up, and the tube colours it.
Rasp rasp	0.0 % ... 100.0 %	0.0 %	Irregular closure of the valve, cycle by cycle. A fox's scream and the break at the top of a howl are both this rather than more harmonics.
Radiate radiate	0.0 % ... 100.0 %	35.0 %	How much of the output is the <i>rate of change</i> of the flow rather than the flow. A small source radiates the derivative, which tilts the harmonics up by 6 dB an octave.
Partials partials	0.0 % ... 100.0 %	60.0 %	How much of the timbre is the archetype's own <i>measured</i> balance between its first six partials rather than the synthetic valve. This is the control that reaches for the recording itself: a fox opens and closes its harmonic stack within one scream, and a fixed valve through a fixed tract cannot do that at all. Scaled by how much of that call's energy the measurement actually accounts for, so an archetype with a second animal in it does not pretend to know.

PHRASE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Calls calls	1 ... 32	3	How many calls one phrase holds. Measured at a median of 3, and the spread is the whole point: a tawny owl's phrase is two notes, a wolf's is one held note, a fox's six and a scops owl's fifty-seven of the same pip.
Call Rate call_rate	0.05 Hz ... 20.0 Hz	0.98 Hz	How fast the calls of one phrase follow each other, onset to onset. Measured across the library at 0.98 Hz – a call about every second – and from 0.32 Hz for a screech owl's spaced whinnies to 2.85 for an owl's paired hoots. A call longer than the interval pushes the next one out rather than being cut short: one animal cannot overlap itself.
Rate Drift rate_drift	0.0 % ... 100.0 %	0.0 %	How much a phrase speeds up as it goes. A screech owl's whinny accelerates; a hoot does not.
Legato legato	0.0 % ... 100.0 %	32.0 %	How far each call is stretched towards filling its own slot, so that a phrase runs together rather than leaving silence between its notes. The default is measured: 32 % of consecutive calls in the library touch, against ChirpParade's 70 % – a night is not a dawn chorus, and a Legato inherited from that plugin stretches a 181 ms hoot to fill an 833 ms slot.
Motif motif	-12.00 st ... 12.00 st	0.00 st	How far the pitch steps from one call of a phrase to the next. A barred owl drops on its last note; most things do not move at all.
Variation variation	0.0 % ... 100.0 %	30.0 %	How much each call differs from the last in contour, pitch and length. At zero a phrase is one call repeated, which nothing in the library is.
Phrase Gap phrase_gap	0.05 s ... 60.0 s	3.70 s	Silence between one phrase and its repeat. Measured at 3.7 s across the library; a tawny owl leaves 6.3 and a scops owl keeps its own up for minutes.
Repeats repeats	1 ... 16	1	How many times a phrase repeats before the animal stops.

PACK

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
-----------	-------	---------	--------------

Shot Level shot_level	-inf dB ... 6.00 dB	0.00 dB	Level of the phrase a note fires. At -60 dB a note brings up the night without anything calling on it.
Pack Level pack_level	-inf dB ... 6.00 dB	-6.00 dB	Level of the animals that call unprompted while a note is held. This is the drone half: at -60 dB every call is one you played.
Pack Rate pack_rate	1.00 /min ... 600 /min	30.0 /min	How often somebody calls, as a Poisson process rather than a clock. Measured 22 to 82 calls a minute across the six callers, median 30.
Animals animals	1 ... 16	3	How many individuals there are. Each has its own pitch, position, distance and voice, held for as long as the note is – a pack is individuals, not one animal moving about.
Pitch Spread pitch_spread	0.00 oct ... 2.00 oct	0.50 oct	How far the pack's pitches spread. Wolves deliberately avoid each other's pitch when they howl together, which is why a chorus of them sounds like more animals than it is.
Voice Spread voice_spread	0.0 % ... 100.0 %	35.0 %	How much the individuals differ in timbre, length, contour and rate.
Answer answer	0.0 % ... 100.0 %	45.0 %	How often one animal answers another, shortly after. This is most of what makes a pack sound like a conversation rather than a random process, and it is what a wolf chorus actually is.
Restless restless	0.0 % ... 100.0 %	40.0 %	A slow drift of how active the pack is. No night is uniform.
Max Voices max_voices	4 ... 64	24	Ceiling on the calls sounding at once, which is the plugin's CPU cost.

CHORUS

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Chorus Level chorus_level	-inf dB ... 6.00 dB	-8.00 dB	Level of the frog chorus. It is a layer of its own because a croak is not a call: it is a pulse train through a body resonance, with no pitch contour in it at all.
Croak croak	0.0 % ... 100.0 %	50.0 %	Which frog. Eight measured croaks, ordered by pulse rate: from a 10 Hz knocking at the bottom to a 69 Hz rattle at the top. Each is the median of one recording – its pulse rate, its pulse count, both of its resonances and its envelope.
Croak Pitch croak_pitch	300 Hz ... 6.00 kHz	2.06 kHz	Where the croak's first resonance sits. Measured 1.5 to 3.3 kHz, median 2063. The second one moves with it: it sits at a measured 0.50 of the first across the library, and a croak with only one resonance is a beep.
Pulse Rate pulse_rate	4.00 Hz ... 200 Hz	27.0 Hz	The rate of the pulse train. This is the species: it is the first thing a field guide lists, and it is what a listener hears as the difference between a knock, a creak and a rattle. Measured 10 to 69 Hz, median 27.
Pulses pulses	1 ... 64	10	How many pulses one croak holds. Measured 2 to 25.
Croak Length croak_length	20.0 ms ... 2.00 s	325 ms	How long one croak lasts. Measured 43 to 608 ms, median 325. It works by scaling how many pulses the croak holds rather than by slowing the train down: a long croak and a slow one are different animals, and the pulse rate is the species.
Frogs frogs	1 ... 48	12	How many frogs there are. Each has its own pitch, position, distance and rhythm, and calls on a period of its own.
Croak Rate croak_rate	2.00 /min ... 600 /min	121 /min	How often one frog croaks. Measured 48 to 480 croaks a minute, median 121.
Regularity regularity	0.0 % ... 100.0 %	95.0 %	How far the pond keeps one rhythm. At zero every frog is independent and the chorus is a Poisson process, which is what the rest of the suite spawns events with – and it is measurably wrong here: the Fano factor of the croak arrivals is 0.90 at 50 ms and 0.30 at four seconds, where a Poisson process is exactly 1.0 at every window. A frog chorus is <i>more</i> regular than random, not less. The default is not a taste judgement either: rendered and measured with the same estimator it gives 0.90 / 0.79 / 0.56 / 0.39 against the library's 0.90 / 0.81 / 0.59 / 0.30, the closest of any setting.

Chorus Spread chorus_spread	0.00 oct ... 2.00 oct	0.70 oct	How far the frogs' pitches spread. A pond is several species at once.
Chorus Width chorus_width	0.0 % ... 100.0 %	85.0 %	How far the chorus spreads across the stereo field.

INSECTS

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Insect Level insect_level	-inf dB ... 6.00 dB	-14.00 dB	Level of the insect bed: crickets and katydids, which are neither callers nor croaks. There is no pitch contour and no body resonance in one – just a narrow band of noise with a trill on it.
Insect Pitch insect_pitch	800 Hz ... 9.00 kHz	2.98 kHz	Where that band sits. Measured at 2.9 to 3.2 kHz in the three references that carry an insect bed clear enough to measure.
Insect Band insect_width	0.0 % ... 100.0 %	35.0 %	How wide the band is. The measured Q is 21, which is narrow – a cricket is very nearly a tone, and widening it turns the bed into a hiss.
Trill Rate trill_rate	2.00 Hz ... 120 Hz	49.2 Hz	The pulse rate on the band. Two references carry a prominent one and they disagree by an octave, 33 and 49 Hz, which is probably a katydid and a cricket rather than an error.
Trill Depth trill_depth	0.0 % ... 100.0 %	70.0 %	How deep that pulsing cuts. All the way down is a chirping insect; not at all is a steady ring, which is what a distant chorus of them becomes.
Shimmer shimmer	0.0 % ... 100.0 %	50.0 %	How much the individuals differ in pitch and trill rate. It is what turns one insect into a field of them.

BED

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Bed Level bed_level	-inf dB ... 6.00 dB	-18.00 dB	Level of the night itself: the wind, the distant water, the chorus too far away to be individual animals. Its shape is measured – the third-octave curve of the quietest third of the frames of all 45 references, solved onto an eight-band filterbank.
Bed Tilt bed_tilt	-100.0 % ... 100.0 %	0.0 %	Tilts that measured curve. Negative is a heavier, further night; positive brings the top back, which is what a still night close to the trees is.
Bed Motion bed_motion	0.0 % ... 100.0 %	30.0 %	A slow drift of the bed's level, so it breathes rather than sitting still.

PLACE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Distance distance	0.0 % ... 100.0 %	45.0 %	How far away the animals are. Distance takes the top off a call and drops its level, and it is the single most effective control here: a wolf at the far end of a valley is most of what the word 'howl' means.
Depth distance_spread	0.0 % ... 100.0 %	45.0 %	How much the animals differ in distance. A night has near animals and far ones, and that spread is most of its depth.
Air air	0.0 % ... 100.0 %	50.0 %	How much of the high end survives the distance. Cold dry air keeps more.
Width width	0.0 % ... 100.0 %	70.0 %	Stereo spread of the pack. They are point sources, so they sit somewhere rather than being everywhere.
Space Amount space_amount	0.0 % ... 100.0 %	35.0 %	How much of the surroundings you hear: a wood, a valley, a marsh.

		%	
Space Size <code>space_size</code>	0.0 % ... 100.0 %	70.0 %	Size of that space, in metres of it. A howl carries because the valley answers it.
Space Damping <code>space_damping</code>	0.0 % ... 100.0 %	55.0 %	How absorbent its surfaces are. A wood full of leaves is very absorbent; a rock face across a lake is not.

FILTER

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Filter Type <code>filter_type</code>	Lowpass, Bandpass, Highpass, Notch	Lowpass	Shape of the output filter.
Highpass <code>highpass</code>	20.0 Hz ... 2.00 kHz	20.0 Hz	Permanent highpass. The lowest fundamental in the library is a 138 Hz howl, so this has room to move before anything is lost.
Filter Cutoff <code>filter_cutoff</code>	30.0 Hz ... 20.00 kHz	20.00 kHz	Corner of the output filter.
Filter Resonance <code>filter_reso</code>	0.0 % ... 100.0 %	10.0 %	Resonance of the output filter.
Filter Key Track <code>filter_keytrack</code>	0.0 % ... 100.0 %	0.0 %	How far the filter follows the played note.

ENVELOPE

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Attack <code>attack</code>	0.50 ms ... 10.00 s	9.76 ms	How long the night takes to come up when a note starts. Longer than ChirpParade's by default: a chorus and a bed fade in, where a single call does not.
Decay <code>decay</code>	5.00 ms ... 20.00 s	400 ms	Fall from the initial level to the sustain.
Sustain <code>sustain</code>	0.0 % ... 100.0 %	100.0 %	Level held while the note is down.
Release <code>release</code>	5.00 ms ... 30.00 s	1.43 s	How long the night takes to stop after the note is let go. Calls already in flight always finish, so a short release thins the pack rather than cutting a howl in half.
Velocity To Level <code>vel_to_level</code>	0.0 % ... 100.0 %	60.0 %	How much velocity sets the level.
Velocity To Pitch <code>vel_to_pitch</code>	0.0 % ... 100.0 %	15.0 %	How much velocity raises the pitch. An animal calling harder pushes more air past a tighter larynx, so it does both.

OUTPUT

PARAMETER	RANGE	DEFAULT	WHAT IT DOES
Output Gain <code>gain</code>	-inf dB ... 12.00 dB	0.00 dB	Output level.
Random Seed <code>seed</code>	0 ... 999	0	Non-zero gives the same night every time. Zero is always different.

The preset library

16 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.

Barred Phrase

Eight notes in a row, which is the other thing an owl does: the barred owl recording segments into eleven calls of a median 240 ms in under four seconds. Rate Drift speeds the phrase up slightly as it goes and Motif walks it down, which is that species' signature and not a general owl behaviour.

Bullfrog Bank

The slow end: a 9.8 Hz pulse rate, which is slow enough to hear the individual pulses rather than a rattle, and four of them to a croak. Croak Pitch is pulled below the measured 2169 Hz of that row, which is the one liberty in this preset – the library has no recording of a true bullfrog, and its lowest measured resonance is 1529 Hz.

Cricket Field

The insect bed on its own, which is the weakest measurement in the plugin and says so. The library has no recording of insects alone: the carrier and the trill rate here come from the cricket band *behind* a scops owl and behind a chorus of frogs, isolated by being narrow, high and steady where everything else in those files is none of the three. 2.9 to 3.2 kHz, pulsed at 49 Hz.

Deep Forest Bed

Nothing calls. This is the measured night on its own – the third-octave curve of the quietest third of the frames of all 45 references, solved onto an eight-band filterbank and then corrected against the engine's own rendered curve until it sat within 3.9 dB rms of it. Bed Tilt pulls it darker, which is what standing under a canopy does.

Fox Across Fields

The same animal at four hundred metres. Nothing about the voice changes except Distance, which is the point: the air absorption and the inverse-distance loss between them turn a scream into the thing people report as somebody shouting in the next field.

Loon Lake

A wail over open water. The loon is the only bird among the callers and the thinnest caller in the table – one recording, eleven usable contours – so its eight archetypes come from less material than any other's. The space is the largest in the factory set at 0.95 with the damping down at 0.3, because what carries a loon is a lake with a rock shore on the far side.

Midnight Marsh

A wet place with something large a long way off. The chorus is the two-pulse creak of `croaking-frogs.wav` at 13 Hz, eighteen of them; the wolf calls about once every fifteen seconds and three animals answer each other at 0.7. Everything is far away – Distance 0.6, Depth 0.55 – which is what makes the space rather than the reverb.

Pack Answering

A chorus of six. Answer is at 0.85, which is the whole preset: one animal calls and another replies from somewhere else a third of a second to two seconds later, and that behaviour – not the voice – is what makes a group of wolves sound like a group. Pitch Spread is wide at 0.85 octaves because wolves in a chorus avoid each other's pitch, which is why a pack sounds like more animals than it holds.

Pond Chorus

Sixteen frogs on one pond. Every number is one row of the croak table – `frogs-6.wav`, measured at a 20 Hz pulse rate, thirteen pulses, a 608 ms croak and resonances at 1535 and 442 Hz – and Chorus Spread puts the pond's other species around it. Regularity is at 0.95, which is not a taste judgement: the measured Fano factor of croak arrivals is 0.30 at four seconds where a Poisson process is 1.0, and this is the setting that renders it.

Scops Metronome

One pip every two and a half seconds, for as long as you hold the note. The scops owl recording is 105 seconds of exactly that – 42 calls at a median 1272 Hz whose measured contour barely moves, a sweep of 0.04 octaves, the flattest thing in the library. Variation is almost off because that is the point of the bird. The insect bed under it is from the same recording, which carries a cricket band at 3.2 kHz.

Screech Whinny

The descending trill, which is why the two screech owls are a caller of their own rather than filed with the hoots: their contours turn direction up to 143 times in a single call, against an owl's ten. The archetypes here are the one part of the library the contour fit under-resolves – 48 terms over a three-second whinny draws 68 of those 143 turns – so Detail is left wide open and the trill is as fine as the table can make it.

Spring Peepers

The high, fast end of the table: `leopard-frogs-croaking.wav`, whose croaks are 43 ms at 3.3 kHz and arrive 480 times a minute – four times the rate of anything else measured. Twenty-four of them at that rate is a wall, and the shortest croak in the library is what makes it a shimmer rather than a rattle.

Summer Night

All four layers at once, which is what the plugin is for: a bed measured from the quiet frames of every reference in the library, a cricket band in it, twelve frogs on a pond at a 48 Hz pulse rate, and an owl that says something every eight seconds or so. The balance is the part worth keeping – the callers are 4 dB under the chorus and the chorus 4 dB over the bed.

Tawny Wood

The two-note hoot, at the measured numbers: the owl group's phrases hold a median of two calls 2.85 Hz apart with 6.3 s between phrases, and its fundamental is 497 Hz with a single harmonic – the purest voice in the library at a measured -34 dB of roughness. Motif drops the second note a tone and a half, which is what the tawny recordings do.

Vixen Scream

Close, and deliberately unpleasant. The fox group is the noisiest thing in the library by a distance: a measured roughness of -23 dB against every other caller's -32 to -34, and three harmonics against their one. Rasp is up at 0.35 because a scream's roughness has structure – the valve does not close the same way twice – and no amount of extra harmonics is the same thing.

Wolf Valley

One wolf, far off. The howl is a measured contour: the eight wolf archetypes run 311 ms to 2.8 s and six of them carry a vibrato of their own at 2.1 to 6.4 Hz, which is what a howl is and what no oscillator with a vibrato on it sounds like. Distance is doing most of the work – the measured group median is 484 Hz and the air absorption at 0.62 takes everything above about 2 kHz off it.