ERIC TRIPPER INSTRUMENTS

WaveFormant

SAMPLE-DRIVEN FORMANT SYNTHESIZER



User Manual

VERSION 1.1 · VST3 · AU

WINDOWS · MACOS

ERICTRIPPERMUSIC.COM

CONTENTS

Welcome

Welcome	2	Header, knobs and shortcuts	12
How WaveFormant works	3	Sounds and your own samples	13
Installation	4	Presets	14
Tour of the instrument	5	Using WaveFormant in a DAW	16
Quick start	6	Sound design recipes	17
The display	7	Performance and troubleshooting	18
Parameter reference	8	Specifications	19

Thank you for choosing WaveFormant

WaveFormant turns any sound into a singing voice. Load a sample — speech, a choir, a flute, a drum loop, the wind — and WaveFormant listens to it. It extracts the formants of the sound, the resonances that give a voice its vowels and an instrument its colour, and sings them back at the pitch you play, with a rich, soft sound source shaped by five resonant formant filters.

Spoken words become melodies. A drum loop starts talking in time. Bells chant, wind turns into a choir. Between a detailed voice that follows every syllable and a slow, gliding pad, WaveFormant covers textures that no other instrument makes from the same sample.

Every control sits on one panel: a live display of the sample and its formants, and eight sections of machined knobs — Scan, Source, Texture, Voice, Morph, Choir, Envelope and FX. Nothing is hidden behind pages, and everything can be automated from your DAW.

What you can do

Sing any sample at any pitch · freeze a single vowel · scan a phrase in loops or ping-pong · morph toward eight vowel and instrument banks · stack up to four choir voices · move from speech detail to pads · blend in the original sample · sync the scan to your song · 72 factory presets.

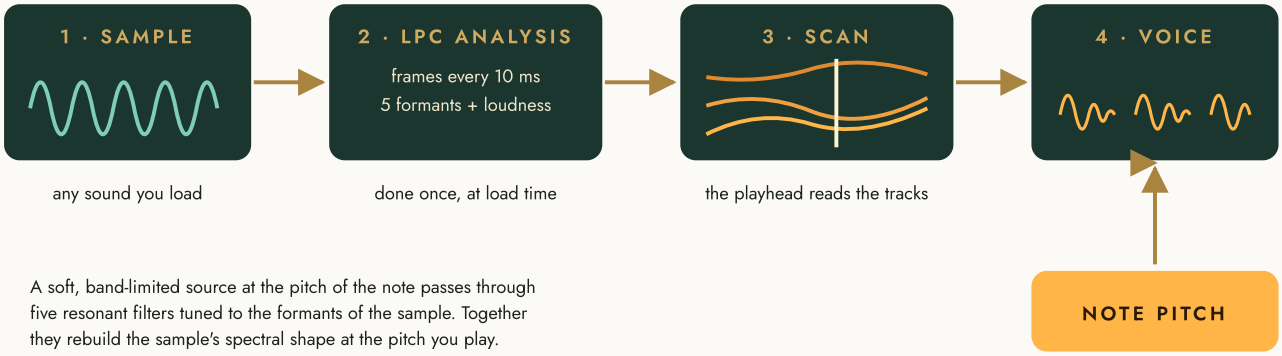
About this manual

Pages 3 to 7 explain the idea and the panel. Pages 8 to 11 describe every setting with its exact range. The rest covers the header and shortcuts, sounds, presets, DAW use, recipes and troubleshooting.

THE IDEA

How WaveFormant works

WaveFormant combines two classic techniques of computer music that are almost never found in plugins: **linear predictive coding** (LPC), which once made computers speak, and the **source–filter model** of the voice, in which a rich sound source is shaped by resonant filters, the way the vocal tract shapes the vocal cords.



A soft, band-limited source at the pitch of the note passes through five resonant filters tuned to the formants of the sample. Together they rebuild the sample's spectral shape at the pitch you play.

1 · ANALYSIS

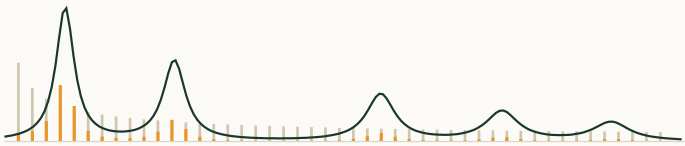
When a sound is loaded, it is mixed to mono, resampled to 16 kHz and cut into 32 ms frames every 10 ms. For each frame, a 14-pole linear-prediction model is computed (pre-emphasis, Hann window, Levinson–Durbin recursion). The peaks of its spectral envelope give up to five formants — frequency, level and bandwidth — plus the loudness of the frame. Quiet frames keep the last voiced formants, and the tracks are gently smoothed. This happens once; the result is drawn on the screen as amber lines over the waveform.

2 · SCAN

A playhead moves through the formant tracks inside a window you define (Position, Window), at any speed — forward, backward, ping-pong, or frozen on a single vowel. It can be shared by all notes (Free) or restart with each note (Retrig), and it can follow your DAW's tempo.

3 · RESYNTHESIS

Each note is a voice made of a **soft, band-limited source** at the pitch of the note — a saw, square or narrow pulse wave, with breath noise and a sine sub-octave on demand — gently low-passed by **Tone**. It runs through **five resonant band-pass filters**, one per formant, whose frequencies, widths and levels follow the formant tracks at the playhead. The filters glide smoothly between frames, without clicks.



grey: source harmonics · dark line: formant filters · amber: what you hear

On top of this, the **Texture** section blends two layers: *Detail* follows every movement of the tracks, *Pad* lets the formants glide slowly into a sustained sound. The **Source** section adds the original sample itself, read from the same scan window and transposed by the keys, and **Mix** balances it against the formant voice.

GETTING SET UP

Installation

REQUIREMENTS

SYSTEM	FORMAT	NOTES
Windows 10 / 11, 64-bit	VST3	Uses Microsoft Edge WebView2 for the interface, already included in Windows 10 and 11.
macOS 10.13 or later, Intel & Apple Silicon	AU · VST3	Universal plug-ins, native on both processors.
Host	—	Any DAW that loads VST3 or Audio Unit instruments: Ableton Live, Logic Pro, Cubase, FL Studio, Reaper, Bitwig, Studio One… WaveFormant is a stereo instrument that receives MIDI notes.

Installing on Windows

1. Unzip **WaveFormant-VST3-Windows.zip**.
2. Copy the **WaveFormant.vst3** folder into `C:\Program Files\Common Files\VST3\`
3. Open your DAW and rescan the plug-ins. WaveFormant appears as an instrument under **Eric Tripper**.
   Ableton Live: Settings › Plug-ins › enable "Use VST3 Plug-in System Folders", then Rescan. FL Studio: Options › Manage plugins › Find plugins.
4. Insert WaveFormant on an instrument track and open its window. The first preset, *Speech to Song*, is loaded.

Installing on macOS

1. Double-click **WaveFormant-macOS.pkg**. It installs the Audio Unit and the VST3 in `/Library/Audio/Plug-Ins/`.
2. The plug-ins are not notarised by Apple, so macOS may block the installer the first time. macOS 14 or earlier: Ctrl-click the .pkg › **Open** › **Open**. macOS 15 or later: System Settings › Privacy & Security › **Open Anyway**.
3. Open your DAW. In Logic Pro, WaveFormant is under **AU Instruments › Eric Tripper**; if it is missing, use Settings › Plug-in Manager › Reset & Rescan Selection.

Updating. Close your DAW, replace the WaveFormant.vst3 folder with the new one (on macOS, run the new .pkg), then reopen the DAW. Your projects keep their settings and sounds.

Window size. The window opens at the largest size that fits your screen (60 to 100 % of 1280 × 848 pixels) and keeps that size in the project.

Tour of the instrument



1. **Preset.** ◀ ▶ step through the presets; click the name to open the preset browser. * means edited.
2. **Sounds.** Factory sounds and file import.
3. **Preview.** Plays the original sound from the scan Position.
4. **Hold.** Latches the notes you play.
5. **Drone.** Holds a note so you can tweak with both hands free.
6. **Sync · Tempo · Tap.** Locks scan and echo to the tempo; shows the tempo; Tap sets it.
7. **Output meter.** Peak level; red near clipping.
8. **Volume.** Output level.
9. **Display.** Waveform, formant tracks and playhead. Click to set the scan Position.
10. **Scan.** Where and how the formant tracks are read.
11. **Source.** The original sample, mixed in.
12. **Texture.** Detail ↔ Pad, Smooth, Tone, Jitter.
13. **Voice.** Source waveform, Breath, Sub Oct, Track, Poly / Mono.
14. **Morph.** Eight banks, Morph, Vowel, Shift, Width.
15. **Choir.** Voices, Detune, Spread, Vibrato.
16. **Envelope.** ADSR with its live curve.
17. **FX.** Glide, Warmth, Echo, Space.
18. **Keyboard.** Click to play; lower on a key is louder.

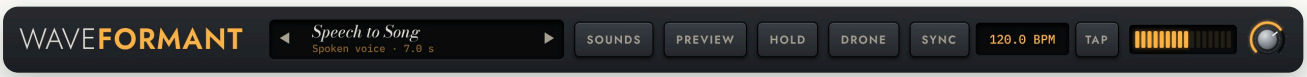
FIVE MINUTES

Quick start

1. **Insert WaveFormant** on an instrument track of your DAW and open its window. The preset *Speech to Song* and the sound *Spoken voice* are loaded.
2. **Play a few notes** on your MIDI keyboard, or click the keyboard at the bottom of the window. The spoken phrase is now sung at the pitch of each note.
3. **Watch the display.** The bright vertical line is the playhead moving through the sample. The amber lines are the five formants found in the sound.
4. In **SCAN**, turn **Speed** to 0: the vowel freezes. Turn **Position**, or click the waveform, to pick another syllable.
5. In **SOURCE**, turn **Mix** toward *Sample*: you hear the original sound, played from the scan position and transposed by the keys. In between, the sample and the formant voice are layered.
6. In **TEXTURE**, turn **Blend** toward *Pad*: the syllables melt into a sustained, gliding sound. Set **Smooth** to taste.
7. In **MORPH**, click **Glass** or **Brass**: Morph jumps to 70 % and the voice takes the colour of the bank. Turn **Vowel** to move through it.
8. Click **Sounds** and load another sound — a choir, a flute, the drum loop — or drag your own audio file onto the window.
9. Use ◀ ▶ to step through the 72 factory presets, or click the preset name to open the browser, try **INSPIRE**, and save your own sound with a name.

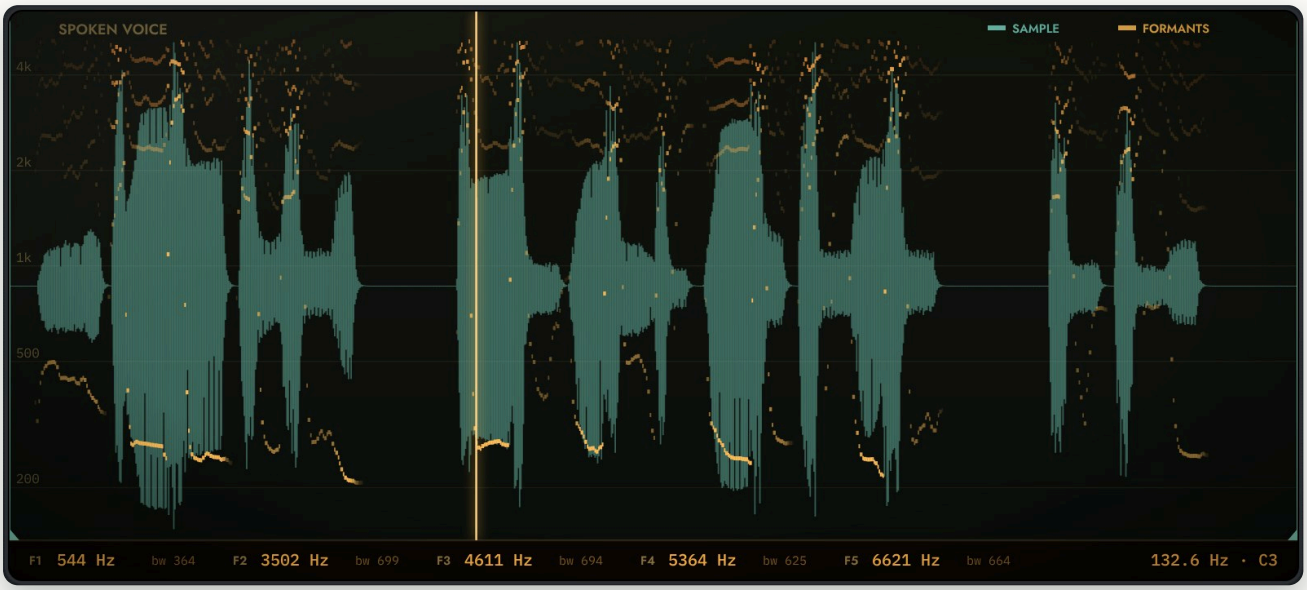
Fine adjustments. Hold **Shift** on your computer keyboard while dragging a knob or turning the mouse wheel. Double-click a knob to return to its default value.

No sound? Check that notes reach the track, and that **Follow** (SCAN) is not reading a silent part of the sample at 100 %.



READING THE DISPLAY

The display



WAVEFORM AND FORMANT TRACKS

The **green-blue** shape is the sample. It is brighter inside the scan window, which is also marked by two small triangles at the bottom. The **amber** dots are the formant tracks found by the analysis, plotted from 150 Hz to 5 kHz on a logarithmic scale (grid at 200, 500, 1k, 2k and 4k); brighter dots are louder formants. The vertical light is the playhead; in Retrig mode there is one per held note. **Click or drag** on the waveform to set the scan Position.

READOUTS

Under the waveform, F1 to F5 show the frequency and bandwidth of the formants currently heard (after Shift, Width, Morph and Track), and the right side shows the pitch of the last note.

SWITCHES

SWITCH	WHERE	WHAT IT DOES
Free · Retrig	SCAN	**Free**: one playhead shared by all notes, running while notes are held. **Retrig**: each note starts its own playhead at Position.
Loop · Ping	SCAN	At the end of the window the playhead jumps back to the start (**Loop**) or reverses direction (**Ping**).
Poly · Mono	VOICE	**Poly**: 8 voices. **Mono**: one voice, legato, last-note priority; use it with Glide.
Soprano … Ghost	MORPH	The bank used by Morph and Vowel. The four vocal banks use measured singing-voice formants; Glass, Brass, Reed and Ghost are instrumental resonance sets. Clicking a bank while Morph is off sets Morph to 70 %, so the change is heard at once; the section title shows the Morph amount.

KNOB COLOURS

Teal arcs belong to the sample (Scan, Source), amber arcs to the voice. An arc turns **orange** when Sync replaces the value by a musical division. Every knob moves on screen when its value changes from outside — mod wheel, MIDI or DAW automation.

Scan & Source

SCAN — Where and how the formant tracks are read.



KNOB	RANGE · DEFAULT	WHAT IT DOES
Position	`0 – 100 % of the sample, shown in seconds · 0`	Start of the scan window. Also set by clicking the waveform.
Window	`0.2 – 100 % of the sample (log), shown in ms / s · 100 %`	Length of the zone the playhead reads in a loop. Very short windows hold a single phoneme.
Speed	`−2× … +2× · 1×`	Playhead speed relative to the original timing. 0 shows **Freeze**: the formants stop on one frame. Negative values read backward. With **Sync** on, this setting becomes a musical division (4 bars … 1/16, triplets and dotted): the time the playhead takes to cross the window.
Follow	`0 – 100 % · 70 %`	How much the loudness of the sample shapes the sound. At 100 % the silences and syllables of the sample gate the voice; at 0 % the voice stays at full level.

SOURCE — The original sample, mixed with the formant voice.



KNOB	RANGE · DEFAULT	WHAT IT DOES
Mix	`Formant … Sample · Formant`	Balance between the formant voice and the **original sample**. At *Sample* you hear the sound you loaded, played from the scan window; in between both are heard (e.g. 70·30).
Tune	`−24 … +24 semitones · 0`	Transposes the original sample.
Key Track	`0 – 100 % · 100 %`	At 100 % the sample follows the keys like a sampler (C4 plays it at its own pitch); at 0 % it keeps its pitch on every note. Glide and vibrato apply.
Loop Fade	`0 – 100 ms · 10 ms`	Crossfade at the loop point of the window, to avoid clicks. Ping mode reverses instead of looping.

Texture & Voice

TEXTURE — The character of the formant voice, from detail to pads.



KNOB	RANGE · DEFAULT	WHAT IT DOES
Blend	`Detail … Pad · Detail`	Mix of the two layers. **Detail** follows every movement of the formant tracks. **Pad** uses smoothed formants and loudness: a sustained, gliding sound. In between both layers are heard (e.g. 50·50).
Smooth	`20 ms – 4 s (log) · 800 ms`	How slowly the Pad layer follows the tracks. Long values give slow, evolving vowels.
Tone	`500 Hz – 16 kHz · 3.4 kHz`	Gentle low-pass on the source before the formant filters. Lower for warm, round voices; higher for bright, airy ones.
Jitter	`0 – 100 % · 15 %`	Slow random pitch drift of each voice (up to ±14 cents): the small instability of real singers and choirs.

VOICE — The sound source, before the formant filters. Poly / Mono sits in the section title.



KNOB	RANGE · DEFAULT	WHAT IT DOES
Source	`Saw › Square › Pulse · Saw`	Waveform of the sound source, band-limited. **Saw**: full and warm, the most vocal. **Square**: hollow, clarinet-like. **Pulse**: narrow and nasal, bright. Intermediate positions blend the neighbours.
Breath	`0 – 100 % · 0`	Adds noise to the source before the filters: from a breathy edge to a pure whisper.
Sub Oct	`0 – 100 % · 0`	Adds a sine wave one octave below the note, for weight and depth.
Track	`0 – 100 % · 0`	Formant key-tracking. At 0 % the formants stay where the sample put them, like a real voice. At 100 % they move with the note (reference C4): useful to keep instruments and bells consistent across the keyboard.

Morph & Choir

MORPH — Reshape the formants and morph them toward a bank.



KNOB	RANGE · DEFAULT	WHAT IT DOES
Morph	`Off – 100 % · Off`	Pulls the formants of the sample toward the bank selected with the buttons above (Soprano … Ghost) and the shape chosen with Vowel. The level stays steady: Morph changes the colour, not the loudness. Clicking a bank while Morph is off sets it to 70 %.
Vowel	`5 shapes per bank (continuous)`	Walks through the five shapes of the bank: a · e · i · o · u for Soprano, Alto, Tenor and Bass; Crystal · Bowl · Prism · Chime · Flint for Glass; Muted · Horn · Bugle · Bell · Blare for Brass; Clarinet · Oboe · Bassoon · Sax · Harmonium for Reed; Hollow · Wraith · Siren · Mist · Echo for Ghost. The display shows the names; in between, two names are shown (e.g. a·e).
Shift	`−12 … +12 semitones · 0`	Moves all formants up or down: bigger or smaller "throat", without changing the pitch.
Width	`×0.25 – ×4 (log) · ×1`	Multiplies the formant bandwidths. Narrow = ringing, whistling resonances; wide = soft, blurred vowels.

CHOIR — Thicken and widen.



KNOB	RANGE · DEFAULT	WHAT IT DOES
Choir	`1 – 4 voices · 1`	Number of voices per note, each with its own pitch drift. With one voice, Detune and Spread have nothing to act on: their knobs are dimmed.
Detune	`0 – 40 cents · 12 cents`	Spread of the choir voices around the note, symmetrically: from a tight unison to a wide, beating ensemble.
Spread	`0 – 100 % · 60 %`	Stereo width: places the choir voices across the stereo field, and alternate formants left and right.
Vibrato	`0 – 100 cents · 20 cents`	Vibrato depth. It fades in after the vibrato delay (0.4 s) at 5.2 Hz; rate and delay can be automated from the DAW.

Envelope & FX

ENVELOPE — Amplitude envelope of each note, drawn above the knobs.



KNOB	RANGE · DEFAULT	WHAT IT DOES
Attack	`1 ms – 5 s (log) · 30 ms`	Rise time of the amplitude envelope.
Decay	`10 ms – 5 s (log) · 400 ms`	Time to fall to the sustain level.
Sustain	`0 – 100 % · 80 %`	Level held while the key is down. 0 % makes plucked sounds.
Release	`10 ms – 8 s (log) · 600 ms`	Fade-out after the key is released.

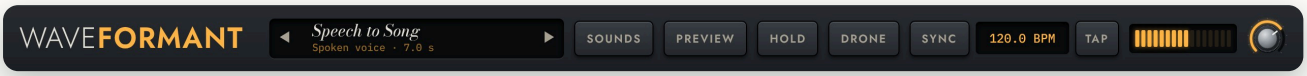
FX — Movement and space.



KNOB	RANGE · DEFAULT	WHAT IT DOES
Glide	`Off – 1 s · Off`	Portamento between notes; most musical in Mono.
Warmth	`0 – 100 % · 25 %`	Soft tape-like saturation of the voice before the effects.
Echo	`0 – 100 % · 20 %`	Stereo echo with a darkened repeat. One control raises both the send and the feedback. The echo time follows the tempo when Sync is on.
Space	`0 – 100 % · 40 %`	Reverb amount and size together: from a small room to a long hall (about 1 to 8 s).

HANDS ON

Header, knobs & shortcuts



THE HEADER

CONTROL	WHAT IT DOES
◀ ▶ and preset name	Step to the previous or next factory preset. Click the name to open the preset browser (factory presets, user presets, Save, Inspire). The sound and its length are shown under the name.
Sounds	Opens the sound browser: 18 factory sounds and file import.
Preview	Plays the original sound from the scan Position, through the effects. Click again to stop.
Hold	Latches every note you play; click again to release the notes that are no longer held.
Drone	Holds a note (the C of the keyboard's current octave) so you can tweak without playing. `Space` does the same.
Sync	Locks the scan speed and the echo time to the tempo. The Speed knob then shows musical divisions and turns orange.
Tempo · Tap	Shows the tempo in use: the DAW's (DAW), a MIDI clock (MIDI) or the internal tempo (BPM). Tap twice or more to set the internal tempo; Tap is greyed out when the DAW or a MIDI clock gives the tempo.
Meter · Volume	Output level, and the output volume knob.

KNOBS

Drag a knob up and down, or use the mouse wheel. Moving slowly makes fine steps, moving fast makes big ones; hold `Shift` for five times finer movements. Double-click returns the knob to its default value. With the keyboard focus on a knob, the arrow keys step the value. The value is shown in the box under the knob.

ON-SCREEN KEYBOARD

The keyboard spans C2 to C7. Click higher on a key for a softer note, lower for a louder one; slide across keys to glide from note to note. The octave buttons set which C the computer keyboard and Drone use.

COMPUTER KEYBOARD

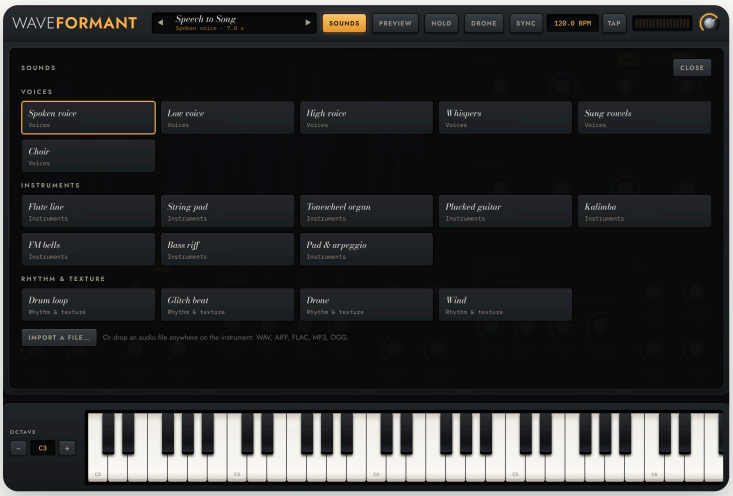
`A` `W` `S` `E` `D` `F` `T` `G` `Y` `H` `U` `J` `K` `O` `L` `P` play notes chromatically from C · `Z` `X` octave down / up · `Space` Drone. In a DAW, some hosts keep the computer keyboard for their own shortcuts: use MIDI or the on-screen keyboard if the keys do not reach WaveFormant.

MIDI

The mod wheel (CC1) moves **Morph** and CC2 moves the scan **Position**; their knobs follow on screen.

SOURCES

Sounds & your own samples



VOICES

Spoken voice, Low voice, High voice, Whispers, Sung vowels, Choir

INSTRUMENTS

Flute line, String pad, Tonewheel organ, Plucked guitar, Kalimba, FM bells, Bass riff, Pad & arpeggio

RHYTHM & TEXTURE

Drum loop, Glitch beat, Drone, Wind

Importing a sound

Drag an audio file from your desktop onto the WaveFormant window, or open **Sounds › IMPORT A FILE...**. WAV, AIFF, FLAC, MP3, OGG and M4A are accepted. Stereo files are mixed to mono for the analysis. The analysis takes a moment; "Analysing formants..." is shown meanwhile.

The loaded sound and its analysis are saved inside your DAW project: the project plays back the same sound on any computer, even if the original file is gone. The audio itself is stored for sounds up to 5 minutes (used by Preview); the formant tracks are always stored.

What makes a good source

Speech & singing

The clearest results. Dry recordings, one voice, no music behind. Poems, counting, radio voices, foreign languages.

Instruments

Solo lines with a strong timbre: flute, oboe, brass, organ, guitar. Raise **Track** so the timbre follows the keyboard.

Rhythm & noise

Loops and textures give rhythmic, talking or windy results. Keep **Follow** high to keep the groove of the loop.

> **Tip.** Reverb and dense mixes blur the formants. Trim long silences before importing, and normalise quiet files: the analysis gates frames that are 30 dB below the loudest part.

STARTING POINTS

Presets



BROWSING

Use ◀ ▶ in the header to step through the 72 factory presets, or click the preset name to see them all. A preset loads its sound with it.

SAVING YOUR SOUNDS

In the browser, type a name and press **SAVE**. A user preset stores every setting and the name of its sound. User presets are shared by all your projects; a preset that uses an imported file needs that sound to be loaded again.

INSPIRE

INSPIRE builds a random but musical patch on the current sound. **Shift + INSPIRE** returns to the initial settings.

An asterisk after the preset name means it has been edited.

PRESET	SOUND	CHARACTER
Speech to Song	Spoken voice	The essential: the spoken phrase sung at the pitch of the keys.
Speech Pad	Spoken voice	The same voice as a slow, wide three-voice pad.
Grain & Nappe	Spoken voice	Both layers at once: articulated detail over a gliding pad.
Choir Wave	Choir	Four-voice choir scanned at half speed into a long, airy wave.
Flute Breath Pad	Flute line	Breathy flute formants as a soft pad.
Organ Mist	Tonewheel organ	Organ resonances smoothed into a misty, tonal bed.
Wind Choir	Wind	The wind sample turned into a whispering choir.
Talking Choir	Spoken voice	Four detuned voices following every syllable.
Robot Psalm	High voice	Square source and sub-octave: a chanting machine.
Whisper Machine	Whispers	Pure breath: rhythmic whispering at any pitch.
Frozen Syllable	Spoken voice	One frozen phoneme, held and widened.
Choir Ghost	Choir	Choir pulled toward the Ghost bank: hollow and eerie.
Vowel Drift	Sung vowels	Sung vowels scanned slowly for evolving colours.
Drum Mouth	Drum loop	The drum loop talks in time, bright source.
Talking Beat	Drum loop	Mono, sub-octave groove driven by the loop.
Bell Throat	FM bells	FM bells with key-tracked, ringing formants.
Bass Formant	Bass riff	Mono bass with a vocal snarl.
Drone Mantra	Drone	Slow scan of the drone: a meditative mantra.
Arp Speaker	Pad & arpeggio	The arpeggio's rhythm and colours sung at your pitch.
Ping-pong Scan	Spoken voice	A short zone of the phrase read back and forth.
Retrig Words	Spoken voice	Each note restarts the phrase: playable words.
Cathedral Speech	Spoken voice	Speech as a huge, reverberant choir pad.
Radio Preacher	Spoken voice	A narrow, warm voice, like an old radio.
Soprano Lines	High voice	Light, high choir with vibrato.
Bass Monk	Low voice	Deep monk choir with sub-octave.
Whisper Cloud	Whispers	Whispers turned into a breathy cloud.
Hushed Pad	Whispers	Soft, warm whispered pad.

Presets (continued)

PRESET	SOUND	CHARACTER
Vowel Organ	Sung vowels	Sung vowels played like an organ.
Sung Glass	Sung vowels	Vowels morphed toward the Glass bank.
Choir in Fog	Choir	A choir drowned in fog: dark and slow.
Brass Chorale	Choir	Choir pulled toward brass resonances.
Child Voice	High voice	Formants shifted up: a small, bright voice.
Giant Voice	Low voice	Formants shifted down with sub-octave.
Throat Singer	Low voice	Narrow formants and sub-octave: overtone singing.
Narrow Vowels	Sung vowels	Very resonant, whistling vowels.
Half & Half	Spoken voice	Half formant voice, half original sample.
Shadow Sampler	Flute line	Mostly the sample, with a formant shadow.
Octave Ghost	Choir	The sample an octave down under a choir.
Fixed Pitch Words	Spoken voice	The sample keeps its pitch; only the voice follows.
Looped Breath	Whispers	A soft breath loop with crossfade.
Strings Doubled	String pad	Strings doubled by their own formant voice.
Flute Choir	Flute line	A choir of breathy flutes.
Talking Strings	String pad	Strings that articulate syllables.
Strings Halo	String pad	A wide, slow string halo.
Organ Vowels	Tonewheel organ	Organ with changing vowels.
Gospel Organ	Tonewheel organ	Warm, chorused organ choir.
Pluck Speaks	Plucked guitar	Plucked notes that speak.
Nylon Mouth	Plucked guitar	Nylon guitar with a vocal mouth.
Kalimba Talk	Kalimba	Kalimba notes with tracked formants.
Music Box Choir	Kalimba	A small music box choir.
Bell Vowels	FM bells	Bells morphed toward Glass.
Bell Pad	FM bells	Bells smoothed into a pad.
Acid Mouth	Bass riff	Aggressive mono bass with glide.
Sub Talker	Bass riff	Very deep mono talking bass.
Arp Choir	Pad & arpeggio	An arpeggio that sings in three voices.
Arp Glass	Pad & arpeggio	Arpeggio morphed toward Glass.
Beatbox Choir	Drum loop	The drum loop as a vocal beatbox choir.
Vowel Snare	Drum loop	Retriggered snare vowels.
Glitch Talker	Glitch beat	Glitches turned into a talking voice.
Glitch Monks	Glitch beat	Glitches sung by low monks.
Stutter Words	Spoken voice	Short retriggered syllables.
Drone Choir	Drone	A bright, sung drone.
Dark Drone	Drone	A dark, warm drone.
Wind Vowels	Wind	The wind sings vowels.
Storm Throat	Wind	A guttural, stormy voice.
Frozen Choir	Choir	One frozen instant of the choir.
Slow Reader	Spoken voice	Speech read in slow motion.
Backwards Prayer	Spoken voice	Speech read backwards.
Ping-pong Vowels	Sung vowels	Back-and-forth through the vowels.
Vowel Morph Pad	Sung vowels	A pad gliding between vowels.
Shifting Voices	Spoken voice	Shifted voices in motion.
Init	Spoken voice	Default settings on the spoken voice.

IN THE STUDIO

Using WaveFormant in a DAW

NOTES

WaveFormant plays MIDI notes from its track: 8 voices in Poly, one legato voice in Mono, velocity-sensitive. All-notes-off messages stop every voice.

TEMPO AND TRANSPORT

When **Sync** is on, WaveFormant uses your project's tempo (the screen shows it with "DAW"). The scan crosses the window in the chosen division and the echo follows the beat. When the transport starts, the playhead returns to Position, so a synced scan starts on the bar every time you press Play.

AUTOMATION

All settings are exposed to the DAW with clear names and can be automated, including a few that have no knob on the panel:

SETTING	RANGE	USE
Internal tempo	`40 – 240 BPM`	Tempo used by Sync when the host gives none.
Sync	`Off · On`	Same as the Sync button.
Speed sync / Echo sync	`4 bars … 1/16`	Divisions used when Sync is on.
Echo time	`40 – 1200 ms`	Echo delay when Sync is off.
Echo tone	`500 Hz – 12 kHz`	Darkness of the echo repeats.
Space damping	`0 – 100 %`	How fast the reverb darkens.
Pre-delay	`0 – 150 ms`	Gap before the reverb starts.
Ensemble	`0 – 100 %`	Three-tap chorus for width.
Tilt	`±5 dB per formant`	Brighter or darker formant balance.
Vibrato rate / delay	`0.5 – 9 Hz · 0 – 2 s`	Speed of the vibrato and the time before it starts.

> **Long parameter list?** Like most VST3 instruments, WaveFormant also exposes technical MIDI-controller entries to the host. They can be ignored; WaveFormant's own settings carry the names shown above.

PROJECTS

Saving the project saves every setting, the loaded sound and its formant analysis, and the name of the preset. The project reopens and plays without opening the WaveFormant window. The window opens at the largest size that fits your screen (60 to 100 %) and keeps that size in the project.

IDEAS

Sound design recipes

Sing any speech

SPOKEN VOICE

Follow 85 %, Blend Detail, Choir 1, Vibrato 30 ct. Play a melody: the words follow your notes.

Frozen vowel pad

ANY VOICE

Speed 0 (Freeze), Window ~1 %, click a vowel on the waveform. Blend Pad, Choir 3–4, Detune 40 %, Space 70 %, slow Attack.

Talking drums

DRUM LOOP

Follow 100 %, Source Square, Tone 80 %, short Release, Sync on, Speed 1 bar. The loop speaks in time with the song.

Whisper choir

WHISPERS OR WIND

Breath 80–100 %, Choir 2–4, Blend 60 %, Space high. Haunting, airy beds.

Ringing glass voice

FM BELLS OR KALIMBA

Track 60–90 %, Width ×0.5, Tone 70 %, Morph 30 % toward Glass. Crystalline, pitched resonances.

Vocal bass

BASS RIFF OR LOW VOICE

Mono, Glide 40 ms, Source Square, Sub Oct 40–60 %, Tone 35 %, Follow 90 %.

Evolving drone

DRONE OR STRING PAD

Speed 0.2×, Ping, Blend Pad, Smooth 2–3 s, Choir 3, Space 75 %.

Playable words

SPOKEN VOICE

Retrig, Window around one word, Speed 1×, Follow 100 %, short Attack: every key says the word.

Ways of thinking

The sample is a mouth, the keyboard is the vocal cords. Everything about "what is said" comes from the sample and the scan; everything about "which note" comes from you. Choose sources for their vowels, consonants and rhythm, not for their pitch.

Detail for words, Pad for colour. Keep Blend near Detail when the articulation matters; push it toward Pad when you want the colour of the sound without its rhythm.

Mix for realism. A little of the original sample (Mix around 20–40 %) under the formant voice keeps the attacks and the texture of the source while the voice sings.

Morph is a tone control. A little Morph toward a vowel or an instrumental bank steadies unruly sources and gives a family resemblance to a set of sounds.

Performance & troubleshooting

PERFORMANCE

The engine is written in C++ and runs inside your DAW like any instrument. Its cost grows with the number of notes and choir voices: a few notes with Choir 1 are light, eight notes with Choir 4 and Blend in the middle are demanding. Rendering or freezing tracks removes the load.

PROBLEM	SOLUTION
WaveFormant does not appear in my DAW	Windows: check that the folder WaveFormant.vst3 is in C:\Program Files\Common Files\VST3 and rescan the plug-ins. Logic Pro: Settings › Plug-in Manager › Reset & Rescan Selection.
macOS says the installer cannot be opened	The plug-ins are not notarised. Ctrl-click the .pkg and choose Open, or on macOS 15 and later allow it in System Settings › Privacy & Security › Open Anyway. The README-macOS.txt file in the download explains the manual installation from the zips.
The window stays blank on Windows	WaveFormant needs Microsoft Edge WebView2. It ships with Windows 10 and 11; on an older or stripped system, install the "WebView2 Runtime" from Microsoft.
I cannot hear my original sound	In SOURCE, turn Mix toward Sample. With Key Track at 100 %, play C4 to hear it at its own pitch.
There is no sound	Check that MIDI notes reach the track. In SCAN, a Follow of 100 % on a silent part of the sample mutes the voice: lower Follow or move Position. Check the Volume knob in the header.
Crackles or high CPU	Each note runs one source per choir voice and ten resonant filters. Lower Choir, keep Blend fully on Detail or Pad (in between both layers are computed), or raise the audio buffer size.
The Morph buttons seem to do nothing	The banks only act when Morph is above 0: the buttons look dimmed while Morph is off. Click a bank (Morph jumps to 70 %) or raise Morph.
An imported file is refused	Use WAV, AIFF, FLAC, MP3, OGG or M4A. Very long files work but take longer to analyse.
An error banner appears in the window	Please send a screenshot of the red banner to the address below: it describes the problem exactly.

REFERENCE

Specifications

Synthesis	Source–filter voice: band-limited source through five resonant formant filters driven by LPC formant tracks analysed from a sample
Analysis	Mono, 16 kHz, 32 ms Hann frames every 10 ms, pre-emphasis, 14-pole LPC (Levinson–Durbin), five formants with level and bandwidth, loudness envelope
Voices	8-voice polyphony or mono legato · up to 4 choir voices per note
Layers	Formant voice (Detail and Pad, blendable · Pad smoothing 20 ms – 4 s) and original-sample layer, balanced with Mix
Sources	Saw · square · narrow pulse (PolyBLEP, continuous blend) · breath noise · sine sub-octave · Tone low-pass
Morph banks	Soprano, Alto, Tenor, Bass, Glass, Brass, Reed, Ghost · 5 shapes each
Effects	Warmth saturation, ensemble, stereo echo with tone control, convolution reverb with pre-delay and damping
Sync	Host tempo and transport; internal tempo and tap tempo
Sounds	18 factory sounds · import WAV, AIFF, FLAC, MP3, OGG, M4A by drag and drop
Presets	72 factory presets · user presets · Inspire
Formats	VST3 · Audio Unit · 64-bit
Systems	Windows 10 / 11 · macOS 10.13 or later (Intel & Apple Silicon)
Interface	One panel, 1280 × 848 at 100 % (60 – 100 %), machined knobs, live formant display
Automation	All settings

Credits

WaveFormant is designed and produced by **Eric Tripper**. More instruments and support: **erictrippermusic.com**.

Linear predictive coding has been used to analyse and resynthesise speech since the late 1960s; the source–filter model behind WaveFormant's voice is the classic description of how the human voice is produced.

TRADEMARKS

VST is a registered trademark of Steinberg Media Technologies GmbH. Windows and WebView2 are trademarks of Microsoft Corporation. macOS, Logic Pro and Audio Units are trademarks of Apple Inc. All other names are the property of their respective owners and are used for identification only.

WaveFormant

ERIC TRIPPER INSTRUMENTS · VERSION 1.1