ORRA AUDIO · PRODUCT MANUAL

Version 1.0.0

ORRA VOCAL

PITCH EDITOR & VOCAL CHAIN

Tune every note of a vocal by hand,
then finish it with eight effects made for the voice,
all in one plugin.

Orra Audio LLC · 2026

Welcome

Orra Vocal is two things in one plugin. **TUNE** is a pitch editor: it shows every note of the vocal on a piano roll and lets you correct, move, retime and reshape each one by hand. Behind it sits a **vocal chain** of eight effects built for the voice: CLEAN, EQ, SAT, COMP, DOUBLE, DELAY, REVERB and OUTPUT.

- **TUNE.** Orra Vocal analyzes the vocal once, in the background, and draws its notes. Tune them all at once with CORRECT, or drag single notes; straighten drift, calm or deepen a vibrato, change a note's level and formant, split, merge and quantize, and shape the slides between notes with GLIDE.
- **A clean sound.** Notes you leave alone play exactly as recorded. Tuned notes keep the singer's tone and character, and the consonants, breaths and silences between them play as recorded.
- **The chain.** Every effect has its own page with large displays and meters. Because Orra Vocal knows where every word, consonant and breath is, CLEAN turns breaths and the noise between phrases down with no threshold to set and no delay.
- **AUTO SETUP.** The first time Orra Vocal hears a vocal, it sets up every effect for that voice, switched off: turn one on and it starts in a sensible place.
- **Ways to judge by ear.** Presets for each effect and for the whole chain, and on every effect page A | B, DELTA (hear only what the effect changes) and MATCH (hear it off at the same loudness).

The vocal reaches TUNE in one of two ways: through **ARA**, in DAWs that support it, or by **TRANSFER**, recording it into the plugin while the song plays. See *ARA and TRANSFER*.

System Requirements

Formats	macOS: VST3, Audio Unit (AU), AAX (Pro Tools). Windows: VST3, AAX. ARA 2 in the VST3 and AU versions.
macOS	10.13 or later (Universal: Apple Silicon and Intel).
Windows	Windows 10 or later, 64-bit.
Channels	Mono, stereo, or mono in and stereo out (a mono vocal with a stereo double, delay and reverb).
Latency	None. The effects add no delay; the analysis runs on its own background thread.
License	A free **14-day demo** with every feature. Buy at **orraaudio.com/products/orra-vocal**; one license activates up to 3 computers. Activation needs an internet connection once.

Getting Started

1. Insert it

Put Orra Vocal on the vocal track as an insert effect, one per vocal track. In a DAW with ARA it reads the vocal clips on that track straight away. In other DAWs, use TRANSFER to record the vocal into it (next chapter).

2. Activate it, or start the demo

The first time you open the window, a panel covers it: *Orra Vocal is not activated on this machine.*

- **Start 14-Day Demo** unlocks every feature for 14 days. Nothing to type.
- **Enter License Key** opens the activation window: paste your license key and click **Activate**.
- **Purchase Orra Vocal** opens the product page.

Until it is activated or the demo is running, Orra Vocal is silent while the DAW plays. When the demo ends the panel comes back with *Your 14-day demo has expired*; enter a license key to carry on.

3. Let it listen

Open the **TUNE** tab. While a vocal is analyzed, TUNE shows **ANALYZING THE VOCAL** with a progress bar; a minute of vocal takes about nine seconds on a recent computer. Then its notes appear on the roll.

4. Tune

For a quick pass, click **CORRECT**: every note is pulled to the nearest note of the **KEY**. For detail, pick a tool on the toolbar and drag notes. Press play in your DAW to hear the result.

5. Switch on the effects you want

Each effect's tab has a light: click it to turn the effect on. AUTO SETUP has already set each one for this vocal, so it starts in a good place. Open the effect's page to fine-tune it.

The license panel

Click the **ORRA VOCAL** wordmark at the top left at any time. The panel shows the version, the license status (with the demo's days left), who it is licensed to, the activations used, and **Deactivate**, which frees this computer's activation so the license can be used on another one. Below are links to the website, support, this manual and Orra Audio's social pages. If the license server can't be reached, Orra Vocal stays activated offline.

ARA and TRANSFER

To tune a vocal, Orra Vocal needs the whole recording, not only the moment that is playing. There are two ways to give it that.

ARA (the direct way)

ARA (Audio Random Access) lets a plugin read the audio clips on its track directly. In DAWs that support ARA 2, such as REAPER, Logic Pro, Studio One and Cubase, insert Orra Vocal on the vocal track the way your DAW adds an ARA plugin: Orra Vocal sees the track's clips, analyzes them, and plays them back corrected.

- Each Orra Vocal shows and edits only **its own track**.
- Edits belong to each copy of a recording: a duplicated track or a copied clip starts with the original's tuning and is then tuned on its own.
- The tuning is saved with the project. A reopened project analyzes the clips once more (the notes appear after a few seconds), then puts every edit back.
- A click on the ruler above the roll moves the DAW's playhead there.

TRANSFER (any DAW)

In a DAW without ARA, or whenever you prefer, record the vocal into Orra Vocal. TRANSFER appears on TUNE's toolbar when there is no ARA.

- **Arm TRANSFER**, then play the song over the vocal. Orra Vocal records what reaches it; while recording you hear the input itself.
- **Stop** or **disarm**: the recording becomes a take and is analyzed.
- From then on, wherever a take covers the song position, its corrected playback plays in place of the input. Elsewhere the input passes through. Where two takes overlap, the newest wins.
- **CLEAR** forgets every take made in this Orra Vocal.

Takes are saved as files in your Music folder, under **Orra Vocal/Transfers**; the project keeps a reference to them and their edits. To open the project on another computer, copy that folder too.

Live input: with no ARA clip or take at the song position, the chain still works on the live input. CLEAN's BREATH and GAP need an analyzed vocal and do nothing on live input.

Interface Tour

The window is a slate-gray faceplate with one page per tab. It can be resized from its corner. Hover over anything for a tooltip.

1. Header

The **ORRA VOCAL** wordmark (the license panel), then the tabs in signal order: **TUNE**, the **chain key** (CHAIN PRESETS), **CLEAN**, **EQ**, **SAT**, **COMP**, **DOUBLE**, **DELAY**, **REVERB** and **OUTPUT**. Each effect tab carries its power light: click the light to switch the effect on or off, click the name to open its page. An effect that is off dims its name. On the right: **SCALE**, the size of everything in the window, and the round **?** key for the tooltips.

2. The TUNE page

A **toolbar** along the top, in symbols (each one's name is in its tooltip): the eight tools, SNAP, KEY, CORRECT and QUANTIZE (each with an arrow for its settings), MERGE, DELETE and RESET, the view keys, LEAD and OTHER TRACKS, UNDO and REDO, and TRANSFER and CLEAR when there is no ARA. Below it the **roll**, the notes on a piano roll with a ruler and scroll bars. Along the bottom the **inspector**, the selected note's values.

3. An effect page

Every effect page is built the same way. The **riser** along the top holds PRESETS, A | B, RESET, DELTA and MATCH. The **screen** in the middle holds the effect's displays: scopes, graphs and meters that move with the vocal. The **cards** along the bottom hold its knobs and switches, grouped by job.

4. Knobs and values

Drag a knob up or down (Shift for fine control). Double-click a knob to return it to its default. Each knob shows its value in its own unit; while you drag, the effect's displays follow.

5. Signal order

vocal in → TUNE (the corrected playback) → INPUT (OUTPUT page) → CLEAN → EQ → SAT → COMP → DOUBLE → DELAY → REVERB → OUTPUT (DE-ESS, RIDE, GAIN, WIDTH, SAFE) → out

Every effect starts switched off except OUTPUT, so a new Orra Vocal plays the vocal untouched until you tune it or turn an effect on.

TUNE: The Note Editor

Each sung note is a shape on the piano roll, at its pitch (up and down) and its time (across). The rows are shaded like the keys of a piano, and each note sits centered on the note it will sound as.

Selecting

- **Click** a note to select it; **Cmd-click** (Ctrl-click on Windows) adds or removes one.
- **Drag across empty space** to select several. **Select all** is in the right-click menu.
- Commands such as CORRECT, QUANTIZE and RESET work on the selected notes, or on all of them when none is selected.

Moving around

Wheel	Scroll. Over the ruler: through the song.
Cmd + wheel, or pinch	Zoom in time (Ctrl + wheel on Windows).
Alt + wheel	Show more or fewer notes up and down (Option on a Mac).
Shift + wheel	Zoom where the mouse is.
Shift-click a note	Zoom in on it. Shift-click empty space zooms back.
Drag down on the ruler	Zoom in in time.
Scroll bars	Drag to scroll; drag either end to zoom; the wheel over a bar zooms; double-click shows everything.
FIT / SEL / FOLLOW	Show every note / zoom to the selection / scroll along while the song plays.

Hear a note

While the song is stopped, a clicked note plays as it will sound, and plays again at its new pitch as you drag it (on each semitone with SNAP on), so you can find the right note before pressing play. A click on empty space stops it. Turn this off in the right-click menu: **Hear a note when it is clicked**.

The right-click menu

Right-click the roll for every tool and command in words: TOOL, SNAP, the NOTES commands (Correct, Quantize, Merge, Split at the playhead, Delete, Reset, Select all) and VIEW (zoom, Follow the playhead, Show the other tracks, This track is the LEAD, Hear a note when it is clicked, Undo, Redo).

TUNE: The Tools

Pick a tool on the toolbar or in the right-click menu; what a drag does depends on it. Hold **Shift** during any drag for fine moves. Every tool works on all the selected notes at once.

Tool	What it does
PITCH	Drag notes up or down. SNAP decides where they land; hold **Alt** (Option) to flip SNAP for one drag. Double-click a note to snap it into the key. **Cmd-drag** (Ctrl-drag) across empty space draws a new note.
TIME	Drag notes left or right, or drag a note's edge to stretch it. A touching neighbor stretches or shrinks to make room. Double-click puts a note back in time.
SPLIT	Click a note to cut it in two at that point, so each part can be moved on its own. **Split at the playhead** is in the right-click menu.
VIBRATO	Drag a note up for more vibrato, down for less; 0 % is dead steady. Only the quick wobble changes; the note's slow drift stays.
DRIFT	Drag a note down to straighten its slow wandering off pitch (a note that sags or creeps), up to exaggerate it. The vibrato stays.
LEVEL	Drag a note up to make it louder, down to make it quieter. All the way down deletes it.
FORMANT	Drag a note up or down to change the voice's character, brighter and smaller or darker and bigger, without changing its pitch. The consonants of the word move with it.
GLIDE	Where two notes touch, drag the handle between them up for a longer, smoother slide, down for a quicker one. Double-click puts it back to AUTO (see *The Inspector & GLIDE*).

Tips

- Tune with CORRECT first, then fix the few notes it got wrong with PITCH.
- A note that sounds out of tune but sits on its line usually drifts: try DRIFT before PITCH.
- Hard, robotic tuning comes from 0 % VIBRATO and short glides; keep some of each for a natural vocal.
- A word whose ending was split off as its own note can be joined back with **MERGE**.

TUNE: Commands, SNAP & KEY

Command	What it does
CORRECT	Tunes the selected notes (all of them when none is selected). The arrow next to it sets how: **PITCH**, how far each note is pulled to the nearest note of the KEY (100 % lands it exactly; default 100 %), **DRIFT**, how much slow wandering is taken away (default 0 %), and **VIBRATO**, how much vibrato and quick wobble is taken away (default 0 %).
QUANTIZE	Moves the notes' starts onto the song's beat grid. The arrow sets the grid (1/4, 1/8, 1/16, 1/32, 1/8 triplet, 1/16 triplet; default 1/16) and **STRENGTH** (below 100 % tightens the timing and keeps some of the feel). Needs the DAW's tempo.
MERGE	Joins the selected neighboring notes into one.
DELETE	Silences the selected notes. RESET › Level brings them back.
RESET	Puts notes back as they were sung: **Everything**, or only their **Pitch**, **Vibrato and drift**, **Level**, **Formant** or **Time**. **Original notes** also undoes every split, merge and new note.
UNDO / REDO	Take back the last change to the notes, or put it back. TUNE keeps its own history, separate from the DAW's undo; it survives closing the window.

SNAP

Where dragged notes land. **OFF**: freely, cent by cent. **SEMI**: on semitones. **KEY**: on the notes of the KEY. Hold Alt (Option) while dragging to flip it for one drag.

KEY

The scale used by CORRECT, SNAP KEY and a double-click with PITCH.

- **Follow the vocal** uses the key Orra Vocal hears in the vocal (named in the menu).
- **Chromatic (all 12 notes)** snaps to the nearest semitone, whatever the key. The default.
- **Major** and **Minor** list every key: pick the song's key when you know it.

Chromatic is the safe choice for a quick CORRECT: a note is never pulled further than half a semitone. With the song's key set, a note sung between two notes of the key goes to the nearer one of those.

TUNE: The Inspector & GLIDE

The inspector along the bottom shows the selected note's values. Drag a value up or down (Shift for fine moves), or double-click it and type. With several notes selected, a change applies to each of them and they keep their differences.

Value	Meaning
NOTE	The note it will sound as. Drag by semitones, or type a note name: A3, C#4, Bb2.
FINE	How far it sits off that note, in cents (a hundredth of a semitone).
VIBRATO	The note's vibrato: 100 % as sung, 0 % dead steady, up to 200 %.
DRIFT	Its slow wandering off pitch: 100 % as sung, 0 % straight, up to 200 %.
LEVEL	Its level change, from +12 dB down to DELETED.
FORMANT	The voice's character, in cents, up to an octave either way.
TIME	How far its start has moved from where it was sung, in milliseconds.
LENGTH	How long it lasts now, in milliseconds. A change stretches the note's end.
GLIDE	How long the pitch takes to slide into this note from the one before, in milliseconds, when the two touch. **AUTO** follows the singer's own slide.

GLIDE

Where two notes touch, the pitch follows the singer's own slide from one to the next (**AUTO**), so corrected notes join without blips or dips. With the GLIDE tool, drag the handle between two notes up for a longer, smoother slide, or down for a quicker one; this takes the robotic step out of hard-tuned notes. A glide set by hand is centered on the sung slide. Double-click the handle to go back to AUTO.

How Orra Vocal tunes

- Notes you don't touch play **exactly as recorded**.
- The consonants, breaths and silences between tuned notes also play as recorded: no swish between words, and hard "t" and "k" sounds stay crisp.
- A tuned note keeps the original's brightness and its loudness, and even big moves (a third, an octave) keep the singer's own voice instead of turning metallic.
- A word's opening scoop stays part of its note, so CORRECT never leaves the start of a word behind.

TUNE: Other Tracks & LEAD

With Orra Vocal on several vocal tracks of a song (a lead, doubles, harmonies), each one can show the others' notes behind its own, so you can line up a double with the lead or check a harmony against it.

OTHER TRACKS

The **OTHER TRACKS** key shows the notes of the song's other tracks with Orra Vocal faintly behind this track's, each in its own light color, with a legend naming them. The arrow beside it opens their menu:

- Show or hide each track by name.
- **Colors**: for each track, one of a set of colors, the DAW's track color, or the default.
- **This track is the LEAD** (the same as the LEAD key).

With no other track to show, the menu says so: *put Orra Vocal on another vocal track of this song*. The other tracks are only shown: you edit each track in its own Orra Vocal.

LEAD

The crown key marks this track as the **lead vocal**. On the other tracks' Orra Vocal its notes are drawn stronger, with a solid outline, and their legend names it LEAD. Mark one lead per song.

Working with doubles

- Open the double's Orra Vocal with the lead shown behind it. Use **TIME** to line up the starts and ends of words with the lead's, and **PITCH** to land each note on the lead's.
- A double that matches the lead too closely can sound like one voice: leave a little of its own timing and vibrato, or let DOUBLE add width instead.
- For a harmony track, the lead's notes show where each harmony note sits against the melody.

The Vocal Chain

Eight effects in a fixed order, each on its own tab. The corrected vocal from TUNE runs through them from left to right; an effect that is off is skipped.

Effect	Its job
CLEAN	Turns down sibilants, breaths and everything between phrases that is not singing.
EQ	An eight-band parametric EQ with dynamic bands, mid / side bands and DRIVE.
SAT	Five saturation bands that lift and color one region of the voice each.
COMP	Evens out the level so quiet and loud words sit together.
DOUBLE	Copies of the vocal, like a double-tracked take, spread across the stereo field.
DELAY	Echoes in time with the song, which step back while the vocal sings.
REVERB	Orra Space's five spaces with their tone controls and a ducker.
OUTPUT	The level from end to end: INPUT, a late DE-ESS, RIDE, GAIN, WIDTH, SAFE, meters.

Power

Click the light on an effect's tab to switch it on or off, from any page. Lit is on; an effect that is off dims its name. The power is a DAW parameter, so it can be automated.

The riser: tools on every effect page

PRESETS	The effect's presets, with arrows to step through them. See *Presets & AUTO SETUP*.
A | B	Two settings of the effect to compare. The first click on B starts it as a copy of A; change it, then click between them. The arrow copies the slot you hear into the other. Both are saved with the session.
RESET	Every control of the effect back to its default (not its power). For a few seconds after a RESET or a preset the key reads **UNDO**: click it to get back what you had.
DELTA	Hear only what the effect changes: what it takes away or adds, without the rest of the vocal.
MATCH	Hear the vocal with the effect off at the loudness the effect gives it, so you judge the sound, not the level. The number shows how much louder (+) or quieter (–) the effect makes the vocal while it sings.

DELTA, MATCH and the bands' SOLO are listening aids: they switch off when you leave the page and are never saved.

Presets & AUTO SETUP

AUTO SETUP

The first time Orra Vocal analyzes a vocal, it measures it (how loud it is sung and how much the level moves, how busy it is, its range, where its sibilants start, how loud its breaths are, its tone) and sets every effect for it, using the song's tempo for DELAY and REVERB. The effects stay **switched off**: turn one on and it is already in a sensible place. Each effect's preset key reads **AUTO**, and AUTO stays first in its list, so you can go back to it. An effect you have already changed keeps your settings; RIDE is left alone.

Effect presets

Click the preset name on the riser for the list: **AUTO**, the factory presets, then **YOUR PRESETS**. The arrows step through them all. A preset sets every control of the effect, but never its power, INPUT's level, OUTPUT's GAIN or RIDE's TARGET. A ***** after the name means you have changed it since.

- **Save as...** stores what you hear as a new preset; **Save**, **Rename** and **Move to the Trash** act on the one loaded.
- **Show the presets folder** opens it in Finder or Explorer.

Effect	Factory presets
CLEAN	Gentle, Bright Vocal, Dark Vocal, Breaths Down, Breaths Out
EQ	Clean Up, Presence, Air, Warm, De-Harsh, Telephone
SAT	Subtle Glue, Chest, Forward, Air, Tape Warmth, Grit
COMP	Gentle Level, Smooth R&B, Upfront Pop, Rap Tight, Parallel
DOUBLE	Tight Double, Subtle Thickness, Wide Stack, Chorus
DELAY	Slap, Eighth Echo, Dotted Ping-Pong, Quarter Throw, Dark Wash
REVERB	Vocal Plate, Small Room, Ballad Chamber, Big Hall, Spring, Ducked Wash
OUTPUT	Gentle Ride, Tight Ride, Late De-Ess, Ride + De-Ess, Safety Ceiling

CHAIN PRESETS

The **chain key** between TUNE and CLEAN saves and loads the whole chain at once: every effect's settings and which effects are on. **Save the chain as...** stores what you hear; pick a chain to load it, and **Undo the chain load** right after if it was not what you wanted. The links on the key light up while a chain is loaded and unchanged. A chain never touches the tuning, INPUT's level, RIDE's TARGET or OUTPUT's GAIN.

Your presets and chains are files in the Orra Audio folder (on a Mac, ~/Library/Orra Audio/Orra Vocal/Presets), shared by every session and every DAW on the computer.

CLEAN

CLEAN takes the sharp edges off a vocal: it turns down sibilants (s, sh, ch), breaths and the noise between phrases. On an analyzed vocal it already knows where each one is, so it never dulls a bright sung note and adds no delay.

Control	Range (default)	What it does
DE-ESS	OFF to 20 dB (10 dB)	How far sibilants are turned down. Only the top end above FREQ is reduced, so the note underneath keeps its body. Works on live input too: then CLEAN listens for them.
FREQ	3 to 10 kHz (5 kHz)	Where the de-essing starts. Lower also catches “sh” and “ch”; higher trims only the hiss of “s”.
BREATH	–24 dB to OFF (–6 dB)	Turns breaths down. Needs an analyzed vocal.
GAP	–40 dB to OFF (OFF)	Turns down everything between phrases that is not part of a word: room tone, hiss, headphone bleed, mouth noises, clicks, sniffs, and breaths that stand alone. It never touches the singing or its consonants and needs no threshold; a breath that runs into a phrase stays with BREATH. Needs an analyzed vocal.

The displays

SCOPE	The last seconds: what came in (pale) and what CLEAN let through (gold), with the de-essing (orange) and the breath reduction (blue) hanging from the top. Wheel: time span; Shift + wheel: height.
VOCAL MAP	Under the scope: where Orra Vocal found sung notes (gold), sibilants (orange) and breaths (blue). On live input the orange row shows what CLEAN hears as sibilance instead.
SPECTRUM	The vocal in (pale) and out (gold). The shaded band above the FREQ line is the only part the de-esser touches; the orange curve is the cut now, the dashed one the most DE-ESS allows. Drag the FREQ line to move it (Shift: fine); double-click it for 5 kHz.
Meters	IN, DE-ESS (reduction now), BREATH (reduction now) and OUT.

Start with GAP around –15 to –20 dB: the room drops away between phrases while the singing is untouched. All the way down makes the gaps silent, which can sound unnatural under a sparse mix.

EQ

An eight-band parametric EQ to shape the vocal's tone, with an analyzer behind it. Any band can also be a dynamic band that moves only when its region gets loud, and work on the mid or the side of the stereo picture only.

Control	Range (default)	What it does
TYPE	7 types	**Bell** boosts or cuts around the frequency; **Low Shelf** and **High Shelf** lift or lower everything below or above it; **Low Cut** and **High Cut** remove the lows or highs; **Notch** cuts a narrow slice; **Tilt** tips the whole balance darker or brighter.
SLOPE	6 to 96 dB/oct (24)	How steep a cut or shelf is. 12 or 24 sounds natural on a voice; 48 and up is surgical.
FREQ	20 Hz to 20 kHz	Where the band works.
GAIN	±30 dB (0 dB)	How much the band boosts or cuts.
Q	0.1 to 30 (1.0)	How narrow the band is: low values are broad and musical, high values narrow and surgical.
MODE	EQ / DYN (EQ)	EQ shapes the tone; DYN moves the band only when its region gets loud (next page).
ST / M / S	(ST)	Where the band works: both channels, only the middle of the stereo picture (the lead voice), or only its sides (doubles, width).
DRIVE	0 to 100 % (0 %)	Console-style saturation for the whole EQ: the louder notes get thicker and denser while quiet parts stay clean. Level-matched, so it adds color, not volume.

The cards

- **BAND**: the band selector 1 to 8, each number in its band's color (double-click an unused one to switch it on), and the selected band's **BAND ON**, **BYPASS** (keep it on the graph but stop it working) and **SOLO** (hear only this band; the others and DRIVE are skipped).
- **BAND n**: the selected band's MODE, ST / M / S, TYPE, SLOPE, FREQ, GAIN and Q.
- **DRIVE**, and **ANALYZER**: the graph's RANGE (±6, 12, 18 or 30 dB) and what it shows: the vocal coming in (green) and going out (blue), only going out, or nothing.

EQ: The Graph, DYN & Stereo

Working on the graph

Each band is a numbered dot with its curve filled in its color; the white line is the sum, the EQ you hear.

Double-click empty space	Add a band there: a Low Shelf below 100 Hz, a High Shelf above 10 kHz, a Bell in between.
Drag a dot	Frequency and gain. Shift: fine. Cmd (Ctrl): Q.
Wheel on a dot	Q. With Alt (Option): gain.
Double-click a dot	Gain back to 0 dB, Q to 1.
Right-click a dot	Mode, Stereo / Mid / Side, Filter Type, Slope, Solo, Bypass, Delete Band.
Wheel on empty space	The graph's range.

Dynamic bands (MODE: DYN)

A DYN band moves only when its region gets loud: a cut for a harsh syllable or a boomy note that leaves the rest of the vocal alone, or a lift that brings a region forward when it is sung out. GAIN becomes **AMOUNT**.

Control	Range (default)	What it does
THRESH	–60 to 0 dB (–18)	The band's level where it starts to move. Above it, it cuts (or lifts) more the louder its region gets, up to AMOUNT.
AMOUNT	±30 dB	The most it moves: below 0 a cut, above 0 a lift.
ATTACK	0.1 ms to 1 s (10 ms)	How fast it reacts. Short catches a harsh syllable; longer lets the start of each word through.
RELEASE	1 ms to 5 s (100 ms)	How fast it lets go. Too short can flutter; too long keeps it working into the next word.
RATIO	1:1 to 20:1 (3:1)	How hard it moves above THRESH: 2:1 gentle, 4:1 firm, 10:1 and above reaches AMOUNT almost at once.

ATTACK, RELEASE and RATIO sit behind the arrow beside THRESH. On the graph, a DYN band's solid curve is where it is now, the dashed one the most it can do.

Mid and side bands

A band set to **M** works only on the middle of the stereo picture, where the lead voice sits; **S** only on the sides, where doubles and width live. The graph then draws two sums: what the middle hears (white) and what the sides hear (lilac). On a mono track there is no side: M works as ST and S does nothing.

SAT

Five saturation bands. Each pushes one region of the voice into a saturator and lifts it as it saturates, a more musical boost than an EQ's gain, as Orra EQ's saturation bands do.

Control	Range (default)	What it does
TYPE	4 types	**TUBE**: warm and full, the strongest lift. **TAPE**: soft and round, it squashes the loud moments. **CONSOLE**: dense, a mixing desk pushed hard. **EDGE**: starts earliest and bites, for presence and grit.
FREQ	20 Hz to 20 kHz	The center of the region the band saturates.
Q	0.1 to 10 (1.0)	How wide the region is: low values a broad part of the voice, high values a narrow slice.
DRIVE	0 to 100 %	How hard the region is pushed: it gets louder and richer, and its loud moments are rounded off.
ST / M / S	(ST)	Both channels, only the middle, or only the sides, as on the EQ.
QUALITY	OFF / 2x / 4x (2x)	Oversampling, so hard-driven high bands stay clean (no harsh aliasing). OFF uses the least CPU; 2x suits most vocals. A change fades SAT out and back in.

The graph

Frequency across, DRIVE up. Each band is a hump in its color, as tall as its DRIVE, with its type under the dot; a ring round the dot glows while the band is adding something. The **white curve** is the boost a quiet passage gets from all the bands, in dB on the right-hand scale; loud passages get a little less. Behind them, the analyzer shows what SAT adds (gold) over the vocal going out (blue), only the vocal, or nothing.

Double-click empty space	Add a band there, its DRIVE where you clicked.
Drag a dot	Frequency and DRIVE. Shift: fine. Cmd (Ctrl): Q.
Wheel on a dot	Q. With Alt (Option): DRIVE.
Double-click a dot	DRIVE 30 %, Q 1.
Right-click a dot	Type, Stereo / Mid / Side, Solo, Bypass, Delete Band.

The BAND card selects bands 1 to 5 and holds BAND ON, BYPASS and SOLO, as on the EQ. The riser's **DELTA** plays only what SAT adds: each band's boost and harmonics, without the dry vocal.

COMP

COMP evens out the vocal's level so quiet words and loud words sit together in the mix. It adds no latency: no lookahead, nothing to delay the track.

Control	Range (default)	What it does
THRESH	–60 to 0 dB (–20)	The level where compression starts. Lower it to turn down more of the vocal; a few dB on the GR meter for the loud words is a good start.
RATIO	1:1 to 20:1 (3:1)	How firmly the vocal is held down above the threshold. 2:1 to 4:1 evens out a performance gently; 8:1 and up holds it almost flat, like a limiter.
ATTACK	0.1 to 200 ms (10 ms)	How fast a loud word is caught. Short times catch every peak but soften the consonants; longer times let the start of each word come through.
RELEASE	5 ms to 2 s (120 ms)	How fast the level comes back after a loud word. Short sounds tighter and louder but can pump; long sounds smoother.
MAKEUP	–12 to +24 dB (+4)	Turns the compressed vocal back up. Set it so switching COMP on and off keeps about the same loudness, or let MATCH show the difference.
MIX	0 to 100 % (100 %)	Blends the compressed vocal with the untouched one (parallel compression): lower keeps the natural dynamics and still adds body.

The displays

SCOPE	The last seconds of the vocal: what went in (pale), what came out (gold), and the gain reduction hanging from the top, with the threshold as a dashed line. Wheel: time window (1 to 16 s); Shift + wheel: height.
CURVE	Input level across, output level up: the curve THRESH, RATIO, the soft knee, MAKEUP and MIX make, against the 1:1 line, with a dot riding it at the level coming in now.
Meters	IN (peak and RMS), GR (gain reduction, hanging from the top) and OUT.

Starting points

- Pop and R&B leads: ratio 3:1 to 4:1, attack 5 to 20 ms, release 80 to 180 ms, 3 to 6 dB of gain reduction on the loud words.
- Rap: a firmer ratio and quicker times keep every word at the front.
- Parallel: a hard setting (8:1, fast attack) with MIX around 40 % adds density without flattening the performance.
- COMP evens out quick changes; for whole lines sung quieter, let OUTPUT's RIDE do the work first.

DOUBLE

DOUBLE thickens the vocal like a double-tracked take: copies of it, each a little late and a little out of tune, drifting slowly and spread across the stereo field. The copies are only ever late, as a real double is, so DOUBLE adds no latency.

Control	Range (default)	What it does
VOICES	1 to 4 (2)	How many copies. 2 is the classic double, one on each side.
DETUNE	0 to 50 cents (14)	How far out of tune the copies may drift. Around 10 sounds like a real second take; higher gets closer to a chorus.
TIME	15 to 50 ms (35)	How late the copies may fall behind; they wander between 10 ms and this. Longer feels looser; past about 40 ms they start to sound like echoes.
MOTION	0 to 100 % (50 %)	How quickly the copies drift. Low: slow, tape-like drift. High: quicker movement, more like a chorus.
WIDTH	0 to 100 % (100 %)	How far apart the copies sit: 100 % hard left and right, 0 % stacked in the middle with the vocal.
MIX	0 to 100 % (50 %)	The balance between the vocal and its copies: 50 % an equal-level double, 100 % the copies only.

The displays

COPIES	Each copy: across, where it sits in the stereo field; up and down, how far out of tune it is now. The pale disc is your vocal; the dashed lines are DETUNE.
TIMING	How late each copy is over the last few seconds, inside the shaded band (10 ms to TIME). Wheel: time window.
OUTPUT IMAGE	The stereo picture leaving DOUBLE, with its correlation: +1 mono, 0 wide, below 0 the copies cancel in mono.
MONO FOLD	What a mono speaker does to the double: its average tone against your vocal (flat is good). RIPPLE is how uneven it is, LEVEL how much quieter.
IN / OUT	The vocal coming in and the doubled vocal going out.

On a mono track Orra Vocal can run mono in and stereo out, so DOUBLE, DELAY and REVERB can still spread the vocal: set the track (or the plugin's output) to stereo in your DAW.

DELAY

Echoes of the vocal in time with the song. The echoes can darken, bounce left and right, and step back while the vocal sings so they bloom in the gaps.

Control	Range (default)	What it does
TIME	1/32 to 1/1 (1/4)	The echo time as a note length, locked to the song's tempo. A dot makes it one and a half times longer, T is a triplet. 1/32 gives a short slapback.
REPEATS	0 to 95 % (30 %)	How much of each echo comes back again. Low: one or two echoes. High: a long trail. It never runs away.
TONE	0 to 100 % (50 %)	Darkens and thins the echoes so they sit behind the vocal, a little more on every repeat. 0 % keeps them full range; 100 % sounds like a telephone.
PING	on / off (off)	Ping-pong: the echoes bounce between left and right. Stereo only.
DUCK	0 to 100 % (50 %)	Turns the echoes down while the vocal sings, so they come up between words and phrases. 100 % hides them completely under the vocal.
MIX	0 to 100 % (25 %)	The balance between the vocal and its echoes.
WIDTH	0 to 100 % (100 %)	How wide the echoes are; 0 % is mono behind the voice.

TIME choices: 1/32, 1/16T, 1/16, 1/8T, 1/16., 1/8, 1/4T, 1/8., 1/4, 1/2T, 1/4., 1/2, 1/2., 1/1.

The displays

ECHOES	One sung note and its echoes on the song's beat grid. Taller is louder, darker is duller (TONE); up is the left side, down the right. The solid part is what you hear while the vocal sings (DUCK).
SCOPE	The last seconds: the vocal (pale), its echoes (gold) and how far DUCK turned them down (from the top). Wheel: time window; Shift + wheel: height.
Meters	IN, DUCK (how far the echoes are turned down now) and OUT.

Starting points

- A slapback (1/32 or 1/16, no repeats, mono) thickens a rock or country vocal without sounding like an echo.
- A dotted eighth with PING and some DUCK fills the gaps of a pop vocal while keeping the words clear.
- For a "throw" on one word, automate MIX or the power for that word only.

REVERB

Orra Space's five reverbs, with its tone controls and a ducker that lets the reverb bloom between phrases.

Control	Range (default)	What it does
TYPE	5 spaces (PLATE)	**ROOM**: small, you hear its walls. **CHAMBER**: dense and warm. **PLATE**: bright and smooth, the classic vocal reverb. **HALL**: big and soft. **SPRING**: metallic and vintage. A change crossfades.
MODE	VINTAGE / MODERN (VINTAGE)	VINTAGE is warmer, slightly driven and darker; MODERN cleaner and brighter. A change crossfades.
DECAY	0.2 to 10 s (1.8 s)	How long the reverb lasts: the time it takes to fade by 60 dB. Each TYPE has its own range; MIN or MAX in the value means the space stops there.
PRE	0 to 250 ms (30 ms)	Pre-delay: a gap before the reverb starts, so each word stays clear.
MOTION	0 to 100 % (25 %)	Slow movement inside the space that keeps long tails smooth; high values add a gentle drift of pitch.
TONE	−100 to +100 % (0)	Darker to the left, brighter to the right: how fast the highs fade in the tail, and a tilt of the reverb's tone.
LO CUT	20 Hz to 2 kHz (150 Hz)	Takes the lows out of the reverb only (the voice keeps them), so it never gets muddy.
HI CUT	1 to 20 kHz (20 kHz)	Takes the highs out of the reverb only, for softer s and t sounds in the tail.
DUCK	0 to 100 % (0 %)	Turns the reverb down while the voice sings, so it blooms in the gaps.
ATTACK	0.1 to 100 ms (5 ms)	How fast DUCK turns the reverb down when the voice comes in.
RELEASE	10 ms to 2 s (180 ms)	How fast the reverb comes back up after the voice stops.
WIDTH	0 to 100 % (100 %)	How wide the reverb is; 0 % puts it straight behind the voice in mono.
MIX	0 to 100 % (20 %)	How much reverb you hear against the dry voice.

The shortest decay each space reaches: ROOM and HALL 0.2 s, CHAMBER 0.33 s, SPRING 0.38 s, PLATE 0.53 s. ROOM and CHAMBER reach 10 s.

REVERB: Displays & Tone

Display	What it shows
SPACE	The reverb of one click with these settings, measured: the gap at the start is PRE, the gold curve the voice's range, the blue one the highs, and the dashed line falls 60 dB at DECAY. The inset shows the first 100 ms: how fast the echoes build into a smooth wash.
TAIL	The last seconds: your vocal (blue) and the reverb alone (gold), with DUCK's turn-down hanging from the top in red. Where a phrase ends, the dashed line shows how DECAY should let the tail fall. Wheel: time window.
TONE	The reverb's tone curve (TONE, LO CUT, HI CUT and three EQ points) over the reverb's spectrum (gold) and your vocal's (pale).
WIDTH	The stereo picture of the reverb alone: a vertical line is mono, a round cloud is wide. The bar is its correlation: +1 mono, 0 wide.
Meters	IN (the dry vocal), DUCK (how far the reverb is turned down now), WET (the reverb on its own as you hear it) and OUT.

Shaping the tone

In the TONE view, LO CUT and HI CUT are points you can drag, and three more points shape the reverb with an EQ of its own: a dip at 300 Hz to clear mud, a lift in the air, whatever the song needs. The voice itself is never touched.

Drag a point	Move it. Shift: fine. Cmd (Ctrl): narrower or wider.
Wheel on a point	Its width.
Double-click a point	Back to flat, or off.

Starting points

- A pop vocal: PLATE, 1.5 to 2 s, 30 to 50 ms PRE, LO CUT around 200 Hz, MIX 15 to 20 %.
- Keep the words clear in a big space: raise PRE and DUCK before lowering MIX.
- A ballad: CHAMBER or HALL, 2.5 to 3.5 s, with DUCK around 20 to 30 % so the tail fills the gaps.
- If the reverb sounds harsh on the s sounds, lower HI CUT or add OUTPUT's DE-ESS, which catches the sibilance the reverb splashes.

OUTPUT: Level

OUTPUT looks after the vocal's level from end to end: how loud it enters the chain, a second de-esser, automatic fader riding, and what leaves. OUTPUT is on by default.

Control	Range (default)	What it does
INPUT	±24 dB (0 dB)	The level going **into** the chain, before CLEAN. Every effect's defaults were tuned on a vocal at –18 dBFS RMS.
SET		Listens to 2 seconds of singing (if the key reads WAIT, play the song) and sets INPUT so the vocal enters at –18 dBFS RMS. GAIN moves the other way, so the vocal leaves as loud as it came in.
POLARITY	on / off	Flips the vocal's polarity at the start of the chain. Try it when the vocal sounds thin together with another mic of the same performance.
RIDE	on / off (off)	Rides the level toward TARGET, word by word, like a hand on the fader. It moves only while the vocal sings, so breaths, gaps and reverb tails are never pushed up.
TARGET	–40 to –6 dB (–18)	The level RIDE brings the vocal to while it sings. RIDE's **SET** puts it where the vocal already sings, so RIDE evens it out without making it louder or quieter overall.
RANGE	0 to 12 dB (6 dB)	The most RIDE moves the vocal, up or down. 6 dB evens out a performance; more also catches whole lines sung much quieter.
SPEED	20 ms to 1 s (100 ms)	How fast RIDE follows: 60 to 100 ms rides word by word; 250 ms and slower evens out only whole phrases.
DE-ESS	OFF to 12 dB (OFF)	A second de-esser at the end of the chain, at CLEAN's FREQ: it catches the sibilance that COMP and SAT bring back up and that the reverb splashes.
GAIN	±12 dB (0 dB)	The output level. It does not change the sound.
WIDTH	0 to 200 % (100 %)	Stereo width of the whole vocal: 0 % mono, above 100 % wider. No effect on a mono track.
SAFE	on / off (off)	A soft ceiling just under 0 dB that catches peaks that would clip, gently. Below –6 dBFS it changes nothing.

AUTO SETUP stages INPUT (and gives the same back on GAIN) the first time it hears a vocal. Presets and chains never change INPUT, GAIN or TARGET: those belong to the recording, not to a sound.

OUTPUT: Meters & Listening

Listening

- **COMPARE** plays the vocal as it came into Orra Vocal (tuned, no effects) at the same loudness as the processed one: is the chain making it better, or only louder?
- **MONO** folds everything to mono, to check the doubles, the width and the reverb on a mono speaker.

Both are listening aids: they switch off when you leave the page and are never saved.

The displays

IN	The vocal entering the chain, after INPUT: peak (light) and RMS (solid). The gold mark is –18 dBFS, where SET puts the vocal's sung RMS.
RIDE	The last seconds of the sung level as RIDE hears it (pale) and where RIDE puts it (gold), against TARGET (dashed). The lane at the bottom is RIDE's gain; red from the top is DE-ESS. Wheel: 4, 8 or 16 seconds.
LOUDNESS	SHORT-TERM (the last 3 s) and INTEGRATED (everything since RESET, pauses left out) in LUFS to BS.1770-4, and the highest TRUE PEAK in dBTP: keep it under –1 for streaming. It keeps counting while the window is closed.
PEAK / RMS	The level leaving Orra Vocal, left and right: the light bar is the sample peak, the solid bar the RMS over 0.3 s, the pin the last peak held; below, the held peak and the RMS in dBFS.
OVER	Lights when that channel's true peak has gone above 0 dBTP since RESET: it would clip after conversion. Turn GAIN down or switch SAFE on.
SAFE	How far the soft ceiling is turning peaks down now. Reads OFF while SAFE is off.
RESET	Starts INTEGRATED, TRUE PEAK and the OVER lamps again.
VECTORSCOPE	The stereo image of everything leaving: a mono vocal is a vertical line, wider sound spreads sideways. The center is –40 dB, the rim 0 dB.
CORRELATION	+1: both sides the same (mono). 0: unrelated (very wide). Below 0: out of phase, which thins out or cancels in mono.
S/M	The side signal's level against the middle, in dB. WIDTH 200 % raises it by 6 dB; 0 % removes it.

A Typical Session

One way through a lead vocal, from raw take to a vocal that sits in the mix. Every step is optional.

1. Bring it in

Insert Orra Vocal on the vocal track. With ARA, open TUNE and wait for the notes; without, arm TRANSFER, play the song through, and stop.

2. Tune

Set **KEY** to the song's key, or leave it on Chromatic. Click **CORRECT** with PITCH at 100 %, then play the song and listen for notes that went to the wrong place: move them with **PITCH**, or put them back with RESET › Pitch. For a softer result, set PITCH lower in CORRECT's settings and run it again.

3. Fix the details

Straighten sagging long notes with **DRIFT**, calm a wide vibrato with **VIBRATO**, tighten timing with **TIME** or QUANTIZE, and smooth a jumpy join with **GLIDE**. Use **LEVEL** on a word that pokes out.

4. Level

On the OUTPUT page, press INPUT's **SET** while the vocal sings, so the chain sees the level its defaults were made for. AUTO SETUP has usually done this already.

5. Clean, shape, control

Switch on **CLEAN** (add some GAP if the room or headphones are audible between phrases), **EQ** (a low cut under the voice, a cut where it is harsh or boxy), **SAT** for presence and warmth, and **COMP** for an even level. On each page, use MATCH to check that the effect sounds better, not just louder.

6. Space

Add **DOUBLE** for width, then **DELAY** and **REVERB**; DUCK on both keeps the words clear and lets the echoes and the tail fill the gaps.

7. Finish

Back on OUTPUT, switch on **RIDE** and press its SET to even out the lines; add the late **DE-ESS** if the chain brought the s sounds back up. Check **COMPARE** and **MONO**, then save the chain as a CHAIN PRESET for the next song.

Settings, Languages & Mouse

SCALE and the window

The **SCALE** box next to the **?** key sets the size of everything in the window, from 100 % to 200 % in steps of 25, menus and pop-ups included. A size the screen cannot hold is grayed out. The window can also be resized from its corner (at least 1024 × 640 at 100 %). Both are saved with the session.

Tooltips & languages

The round **?** key turns the hover tooltips on and off (on by default). **Right-click** it for the tooltip language: English, Dutch, Japanese, Spanish, Brazilian Portuguese, German or French. Control names such as CORRECT or DE-ESS stay in English, so the tips match the panel.

Saved with your session

Every effect setting and power, A and B, the tuning (every edit to every note), TRANSFER's takes (as references to their files), the tool, SNAP, KEY, CORRECT and QUANTIZE settings, TUNE's zoom and scroll, each page's view settings, and SCALE. TUNE's selection and undo history survive closing the window.

Your presets and chains are yours rather than the project's: every Orra Vocal on the computer shares them. DELTA, MATCH, SOLO, COMPARE and MONO are never saved.

Keyboard & mouse

Your DAW decides whether a plugin receives key presses, so every command in Orra Vocal is on the toolbar, a card or a right-click menu. These mouse moves work in every DAW (on Windows, Ctrl for Cmd; on a Mac, Option for Alt):

Shift-drag (a knob, a note, a dot)	Fine control.
Double-click a knob	Back to its default.
Cmd-click a note	Add it to the selection, or take it out.
Cmd-drag on empty roll (PITCH)	Draw a new note.
Alt-drag a note (PITCH)	Flip SNAP for this drag.
Shift-click a note / empty space	Zoom in on it / zoom back.
Cmd + wheel, Alt + wheel, Shift + wheel	Zoom in time / in pitch / where the mouse is.
Cmd-drag a dot (EQ, SAT, REVERB TONE)	Its Q or width.
Wheel / Shift + wheel on a scope	Its time window / its height.
Right-click the roll, a dot, the ? key	Every command / the band's menu / the language.

Good to Know

A reopened project analyzes again

Orra Vocal does not store the analysis in the project, so a reopened project analyzes its vocals once more (about nine seconds per minute of vocal) and then puts every edit back. The notes appear when it is done.

TRANSFER takes live in your Music folder

A take is too big to keep inside the project, so it is saved under Music/Orra Vocal/Transfers. Moving a project to another computer: copy that folder too. CLEAR forgets the takes but leaves the files in place.

Edits belong to each copy of a recording

A copied clip starts with the original's tuning and is then edited on its own. Two clips the DAW treats as one shared copy of the same audio share their edits: that is the DAW's decision.

QUANTIZE and DELAY follow one tempo

They use the song's tempo as the DAW reports it. In a song with tempo changes, check quantized notes after a change.

GLIDE needs touching notes

A glide is offered only where two notes touch. Between notes with a consonant or a breath in between, the gap plays as recorded.

TUNE's undo is its own

UNDO and REDO on TUNE's toolbar take back note edits; the DAW's own undo does not. On the effect pages, RESET's UNDO and A | B are the quick way back to an earlier setting.

BREATH and GAP need an analyzed vocal

They work from the vocal map that the analysis makes. On live input (no ARA clip, no TRANSFER take) they do nothing; DE-ESS then listens for sibilants by itself.

Mono tracks

On a mono track with a mono output there is no side: PING, the WIDTH controls and S bands do nothing. Give the track a stereo output to let DOUBLE, DELAY and REVERB spread.

Saving CPU

An effect that is off does no work. SAT's QUALITY OFF uses the least CPU, and the displays only draw while the window is open.

Credits & Support

Pitch detection uses the **CREPE** pitch network (J. W. Kim, J. Salamon, P. Li and J. P. Bello, “CREPE: A Convolutional Representation for Pitch Estimation”, ICASSP 2018), its published weights run in Orra Audio’s own code. CREPE: Copyright © 2018 Jong Wook Kim, MIT License. Pitch and formant changes use **Signalsmith Stretch** 1.4.0, Copyright © 2022 Geraint Luff / Signalsmith Audio Ltd., and **Signalsmith Linear** 0.6.4, Copyright © 2025 Signalsmith Audio, both under the MIT License:

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

ARA support uses the **ARA Audio Random Access SDK** by Celemony Software GmbH, under the Apache License 2.0 (apache.org/licenses/LICENSE-2.0), and parts of JUCE’s ARA plugin example, Copyright © Raw Material Software Limited, under the ISC License.

Trademarks

ARA Audio Random Access is a trademark of Celemony Software GmbH. VST is a trademark of Steinberg Media Technologies GmbH. Audio Unit, Logic Pro and macOS are trademarks of Apple Inc. AAX and Pro Tools are trademarks of Avid Technology, Inc. All other names belong to their owners.

Orra Vocal is a product of Orra Audio LLC. © 2026 Orra Audio LLC.

For support, updates, and documentation: orraaudio.com