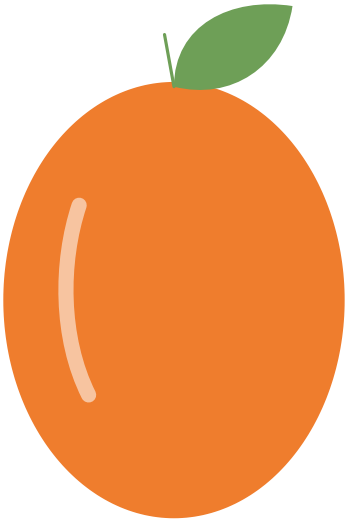


Kumquat Audio

Instruction manual

Edition 1.7.11 · October 2026

kumquataudio.com

BEFORE YOU BEGIN

To the owner

Thank you for choosing Kumquat Audio. These plug-ins are made to be simple: a few controls, and a picture that shows you what is happening.

Each one has a window with a few controls and a picture. Nothing is hidden behind menus. If you want to know what something does, rest the pointer on it: a short note appears. (If it does not, press the “?” at the top of the window to turn the notes on.)

This manual shows each plug-in as a drawing, with every part numbered, and then tells you how to use it. Look at the drawing first, find the part you want, and read its number in the list underneath.

Turn each plug-in on, listen, then turn it off again and listen again. The best setting is the one you can hardly hear working.

How this manual is arranged

The first pages tell you how to install the plug-ins. Then each plug-in has its own chapter, in three parts: **the drawing** with its numbered parts, **the operation** step by step with a table of every control, and **how it works** with the specifications. At the back are the notes on whole-recording plug-ins, care and troubleshooting, a glossary and an index.

Signs used

TIP A tip is something that will help you get a better result.

CAUTION A caution is something that can spoil a recording if you are not careful.

CONTENTS

Contents

	To the owner	2
	Choosing a plug-in or instrument	4
	Installing the plug-ins	5
	In, Out and Gain Match	6

THE PLUG-INS

No. 1	**KA De-Esser**	8
No. 2	**KA De-Click**	12
No. 3	**KA De-Noise**	16
No. 4	**KA De-Verb**	21
No. 5	**KA EQ**	25
No. 6	**KA Limiter**	32
No. 7	**KA Mic Blend**	38
No. 8	**KA Drum Gate**	42
No. 9	**KA Silence Trim**	46
No. 10	**KA Auto Mixer**	50
No. 11	**KA Saturator**	54
No. 12	**KA Delay**	58
No. 13	**KA Bus Comp**	62
No. 14	**KA Reverb**	66
No. 15	**KA Time Align**	70
No. 16	**KA Multiband Comp**	75
No. 17	**KA FET Comp**	79
No. 18	**KA Opto Comp**	83
No. 19	**KA Rider**	87

THE INSTRUMENTS

No. 20	**KA Chladni**	91
No. 21	**KA Root**	95

THE PLUG-INS

CHOOSING

Choosing a plug-in

Find what you need, then turn to the page. **Plug-ins** work on the sound of a track, like effects. **Instruments** make sound from the notes you play, and go on an instrument track.

Number	Name	Fixes	Page
PLUG-INS			
No. 1	**KA De-Esser**	Harsh esses	8
No. 2	**KA De-Click**	Mouth clicks	12
No. 3	**KA De-Noise**	Hiss and hum	16
No. 4	**KA De-Verb**	Too much room	21
No. 5	**KA EQ**	Tone and clashes	25
No. 6	**KA Limiter**	Loudness	32
No. 7	**KA Mic Blend**	Two mics	38
No. 8	**KA Drum Gate**	Drum bleed	42
No. 9	**KA Silence Trim**	Dead air	46
No. 10	**KA Auto Mixer**	Open mics	50
No. 11	**KA Saturator**	A thin sound	54
No. 12	**KA Delay**	Echoes and repeats	58
No. 13	**KA Bus Comp**	A loose mix	62
No. 14	**KA Reverb**	A dry sound	66
No. 15	**KA Time Align**	Mics out of time	70
No. 16	**KA Multiband Comp**	One part of the spectrum out of control	75
No. 17	**KA FET Comp**	Drums and vocals that need bite	79
No. 18	**KA Opto Comp**	Uneven levels	83
No. 19	**KA Rider**	A voice that is too quiet, then too loud	87
INSTRUMENTS			
No. 20	**KA Chladni**	Making music from a sound	91

GETTING STARTED

Installing the plug-ins

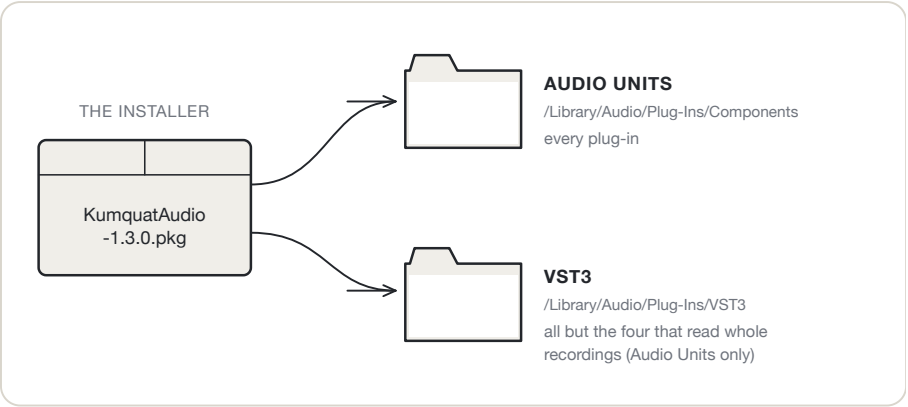


FIG. 1 One installer puts the plug-ins where your music software looks.

Steps

1. Open the installer, KumquatAudio-1.7.0.pkg. It is signed and checked by Apple, so your Mac will not warn you.
2. Choose what to install: Audio Units, VST3, or both. Both are chosen to begin with.
3. Restart your music software. Logic Pro looks at new plug-ins the first time it starts; this can take a minute.
4. Look for Kumquat Audio in your list of effects. The plug-ins are named “KA” and what they do: KA Limiter, KA De-Esser, and so on. The two instruments, KA Chladni and KA Root, are not effects: put them in the instrument slot of an instrument track.

Where to find them

In Logic Pro, open the Audio FX slot of a track, choose **Audio Units**, then **Kumquat Audio**. Other music software lists them under Kumquat Audio as well, in its Audio Unit or VST3 list. To remove them, delete the Kumquat plug-ins from the two folders above and restart your music software.

EVERY INSERT PLUG-IN

In, Out and Gain Match

Along the bottom of most windows is a strip with two knobs, one under each level meter, and two buttons. In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Double-click either knob to return it to 0 dB. Bypass plays the sound without the plug-in (it stays in time with the other tracks, and fades in and out without a click); your music software's own bypass button for the plug-in does the same. The four plug-ins that read whole recordings, and KA Limiter, have the Bypass button at the top of the window instead.

A louder sound almost always seems better, even when it is not. That is how a setting can fool you: you turn something up, the sound gets louder, and you think it got better. Match takes that away. When it is on, the plug-in listens to how loud the sound is as it arrives and as it comes out, and turns the output up or down by the difference, slowly, so the two stay the same. That includes In: drive the plug-in harder and the output stays where it was. Now when you change a setting you hear what it did to the sound, and nothing else. Out is then yours to trim on top.

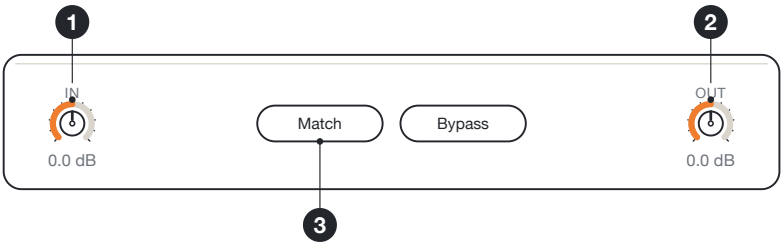


FIG. 2 The strip along the bottom of a window.

1 **In**
Turns the sound up or down before the plug-in hears it.

2 **Out**
Turns what comes out up or down.

3 **Match**
Keeps the output as loud as the sound that came in, whatever In and the plug-in do. The number beside it shows how much it is adding or taking away.

Use In to drive a plug-in harder or softer: more In into KA Saturator or KA Bus Comp gives more effect. Use Out to set the level you want to leave with; with Match on it moves the level up or down from the level that arrived. The number beside Match shows how much it is adding or taking away.

Who has what

In, Out and Match

Plug-in	The strip has	Note
KA De-Esser	In · Out · Match	Match is on hold while Listen is on.
KA De-Noise	In · Out · Match	Match is on hold while Listen is on.
KA De-Verb	In · Out · Match	Match is on hold while Listen is on.
KA De-Click	In · Out	It changes the level so little that matching is not needed.
KA EQ	In · Out · Match	Out is the level after the EQ and the suppressor.
KA Bus Comp	In · Out · Match	Makeup is still there as the compressor's own make-up gain.
KA Saturator	In · Out · Match	Drive already keeps the level close; Match finishes the job.
KA Zing	In · Out · Match	Match is on hold while Delta is on.
KA Contour	In · Out · Match	Match is on hold while Listen is on.
KA Mic Blend	In · Out	In is applied to both mics, including the sidechain.
KA Delay	In · Out	Mix is meant to change the level, so there is no Match.
KA Multiband Comp	In · Out · Match	Each band has its own Gain as well.
KA FET Comp	In · Out · Match	In is the compressor's Input: the threshold is fixed.
KA Opto Comp	In · Out · Match	Out is the compressor's make-up gain.
KA Rider	In · Out	Match would undo the ride, so there is none. Delta plays only the ride.
KA Chladni, KA Root	none	Instruments have no audio input, so no In, Out or Match: use the Level drawing.
KA Reverb	In · Out	Mix is meant to change the level, so there is no Match.

TIP To set a compressor or a saturator fairly: turn Match on, switch the plug-in off and on, and choose the setting that sounds best at the same loudness.

CAUTION KA Limiter has no In and Out: its Threshold is its input and its Ceiling is its output. The four plug-ins that read whole recordings do not have the strip. Every plug-in has a Bypass button.

NO. 1

KA De-Esser

Takes the harsh edge off “s” and “sh”.

KA De-Esser turns down the harsh, whistling part of “s” and “sh” sounds in a voice. It shows you where the esses are, finds them by itself, and can play back only what it removes.

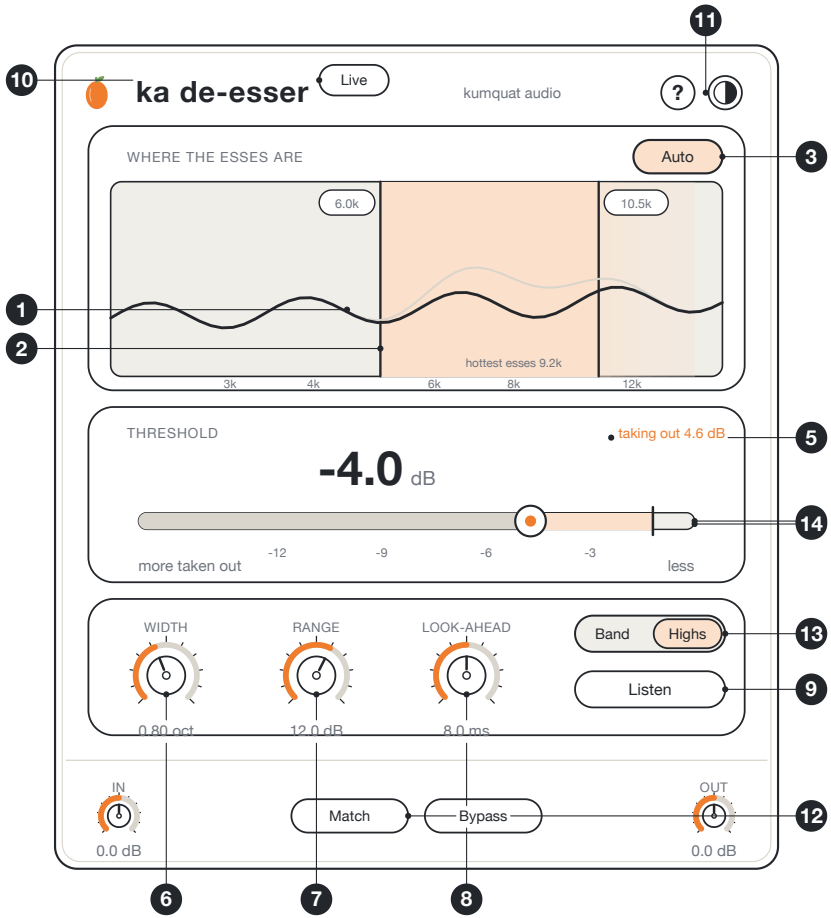


FIG. 3 The parts of KA De-Esser. The numbers are explained on the next page.

Parts

fig. 3

1 The spectrum picture
The sound of the voice. The grey line is what comes in, the dark line what goes out. The orange glow behind them is where this voice's esses stand out (the hottest place is marked): it builds up as the voice talks, so it stays put.

2 The band edges
In Auto they follow the esses, to the middle of the glow. Drag an edge to move it, or the middle of the band to take over by hand.

3 Auto
On: the band finds the esses by itself. Double-click the picture to return to Auto.

4 Threshold
Where the line between a voice and an ess is. Drag the thumb: left takes out more (even a mild ess), right takes out less; at the far right it is off.

5 Reduction
How many dB it is turning the esses down right now.

14 Live meter
The track of the Threshold is also a meter: it fills as the sound gets more essy, and the thin mark is the highest a moment ago. Where the fill goes past the thumb it glows: that part is being turned down. Set the thumb so that the esses go past it and the rest of the voice does not.

6 Width
How wide the automatic band is. Dimmed when you set the band by hand.

7 Range
The most it will turn the esses down.

8 Look-ahead
How far ahead of the ess the reduction starts. The sound is delayed by this much, and the plug-in tells your music software.

9 Listen
Plays only what is being taken out.

13 Band / Highs
Highs turns down everything above the band's lower edge while an ess sounds, as a smooth shelf: the voice below is not touched. Band turns down only the band, a gentler bell.

10 Live
Switches the look-ahead off, so the plug-in has no delay at all: for playing through it.

11 Help and light/dark
The ? turns the pop-up help on and off. The half circle switches light and dark.

12 In, Out and Match
In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

Operation

step by step

1 Put the plug-in on the voice track and play a passage with esses in it.

2 Leave Auto on. As the voice talks, an orange glow builds up in the picture where its esses stand out: the hotter, the harsher. It brightens with every ess, so you can see where each one is. The hottest place is marked, and the band settles in the middle of it. It moves slowly: a person's esses are in much the same place all the time.

3 Watch the Threshold slider: its track fills as the sound gets more essy, and where

Controls and settings

range · start

Control	Range	Starts at	What it does
Threshold	–14 to 0 dB	–4 dB	Where the line between a voice and an ess is: how loud the band may get next to the whole sound before it is turned down. Lower takes out more, even a mild ess; at 0 it is off.
Mode	Band / Highs	Highs	What is turned down while an ess sounds: everything above the band's lower edge, as a smooth shelf, with the voice below it not touched (Highs, the most effective), or only the band itself (Band, a gentler bell).
Look-ahead	0 to 15 ms	8 ms	How far ahead of the ess the reduction starts. The sound is delayed by this much, and the plug-in tells your music software.
Live	on / off	off	Switches the look-ahead off: no delay at all, for playing through the plug-in.
Width	0.3 to 2.0 oct	0.8 oct	How wide the automatic band is. Used in Auto only.
Range	0 to 20 dB	12 dB	The most it will turn the esses down.
Auto	on / off	on	The band follows the esses by itself.
Band edges	2–12 kHz and 3–16 kHz	5 and 10 kHz	Set by dragging the edges in the picture.
Listen	on / off	off	Plays the band: in Auto, only the "s" sounds; set by hand, everything in the band.
In / Out	–24 to +24 dB	0 dB	The two knobs under the meters: before and after the plug-in.
Match	on / off	off	Keeps the output as loud as the input. See In, Out and Gain Match.

TIP Start at –4 dB and lower the Threshold a little at a time. If the voice begins to lisp, you have gone too far.

TIP If the sound you want to catch is very narrow, make Width smaller: it will then catch only the sharpest part.

TIP In Auto the band needs a few esses to settle. In the first seconds of a new voice it sits near 8.5 kHz, where most esses are.

How it works

a closer look

KA De-Esser watches two things all the time: the band, a bell-shaped band-pass round the part of the spectrum where the esses are, and the whole sound. An ess is the moment the band is nearly as loud as everything else put together; a vowel never is.

When the band gets closer to the whole sound than the Threshold allows, the plug-in turns it down, by as much as it is over, but never by more than Range. The sound is delayed by the Look-ahead and the reduction is not, so it begins before the ess does. It lets go in about thirty milliseconds, so it is gone before the next word.

The band's share is taken out of the sound, and the rest of the voice is not touched. In Auto the plug-in keeps two average spectra, of the frames with an ess in them and of all the others, and the band goes to where the esses stand out above the rest, slowly, so that it does not jump about. Width sets how wide it is. The turning down is 1.5 dB for every dB the band is over the Threshold, up to Range, and it is a high shelf from the band's lower edge (Highs) or a bell (Band).

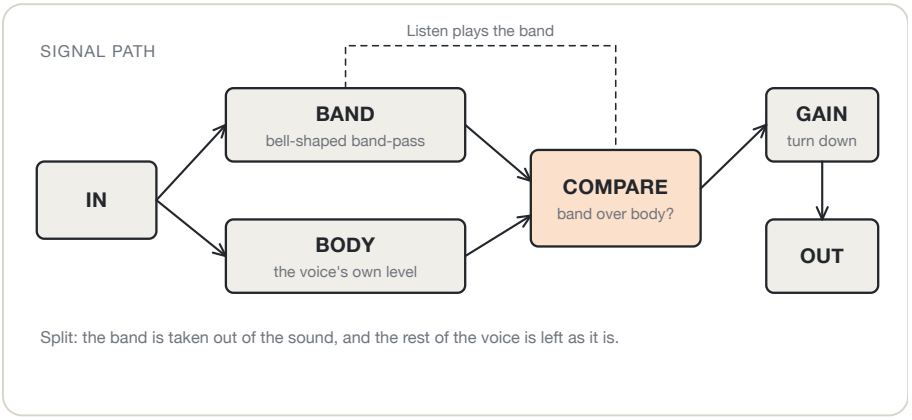


FIG. 4 How it decides: the band is measured against the whole sound.

NO. 2

KA De-Click

Removes lip smacks, tongue clicks and saliva sounds.

KA De-Click removes mouth sounds from a voice: lip smacks, tongue clicks, saliva. A click is a very short burst of high frequencies that jumps out of the sound and is gone again; the plug-in finds it and fades it out for those few milliseconds, and leaves the voice underneath untouched.

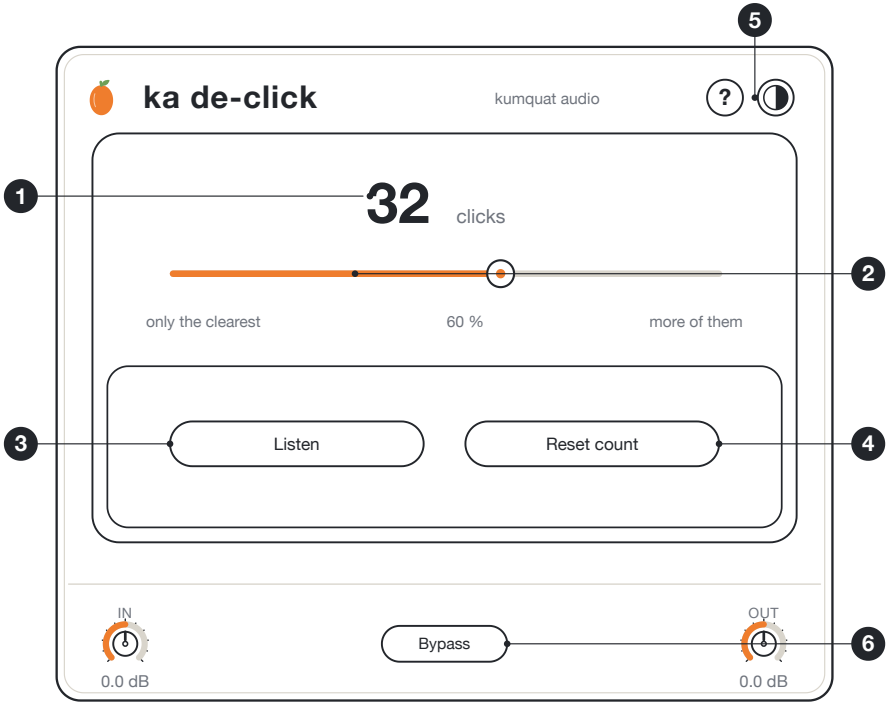


FIG. 5 The parts of KA De-Click. The numbers are explained on the next page.

Parts

fig. 5

1. **Click counter**
   How many clicks have been taken out since you last reset the count.
2. **Sensitivity**
   Left: only the clearest clicks. Right: more of them, and taken out more completely. 0 is off.
3. **Listen**
   Plays only what is being taken out. You should hear clicks and nothing else.
4. **Reset count**
   Starts the click counter again from zero.
5. **Help and light/dark**
   The ? turns the pop-up help on and off. The half circle switches light and dark.
6. **In and Out**
   In turns the sound up or down before the plug-in hears it, and Out turns the result up or down.

Operation

step by step

1. Put the plug-in on the voice track and play a stretch where you can hear the clicks.
2. Turn Sensitivity up until the clicks are gone. The counter shows how many have been taken out.
3. Press Listen. You should hear only clicks. If you hear words or tone, you have gone too far (the sounds of a hard “t” or “k” begin like a click).
4. Press Reset count to start counting from zero again.

Controls and settings

range · start

Control	Range	Starts at	What it does
Sensitivity	0 to 100 %	50 %	How many clicks it finds, and how completely it removes them. 0 is off.
Listen	on / off	off	Plays only what is being taken out.
Reset count	button		Starts the click counter again from zero.
In / Out	In ±24, Out ±12 dB	0 dB	The two knobs under the meters: before and after the plug-in.

CAUTION Very high sensitivity can touch the first part of hard consonants. Always check with Listen.

TIP It finds about twice as many clicks as it used to. Start at 50 % and listen to what is taken with Listen; raise it for a mouth that clicks a lot.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo
Delay	1152 samples (24 ms at 48 kHz)
Window	700 × 304, can be resized

How it works

a closer look

A mouth click is a very short burst of high frequencies. The plug-in splits off the highs with filters that do not shift the phase of the sound, and measures how much sound there is above 6 kHz in every millisecond.

Each millisecond is compared with the median of the milliseconds round it: a few before and a few after. A millisecond that stands well above its surroundings is a click starting. The more you turn Sensitivity up, the less it has to stand out.

A burst that goes on for more than a few milliseconds is speech, and so is a burst after which the highs stay up (the start of a “t” or an “s”): those are left alone. A click is faded out, in the highs only, for those few milliseconds, and everything below it stays, so the voice under the click is kept. To do all this the plug-in looks 24 ms ahead.

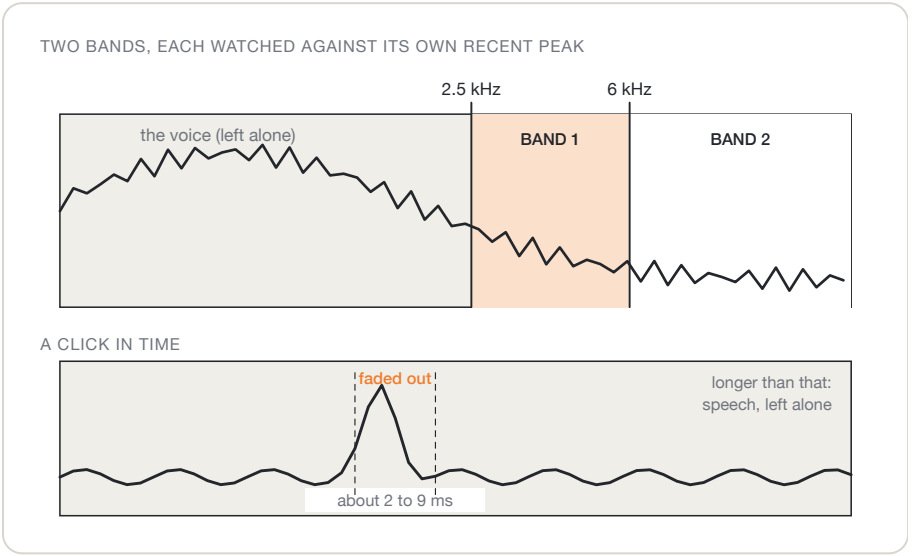


FIG. 6 Two bands are watched, each against its own recent peak.

NO. 3

KA De-Noise

Takes out steady noise: hiss, hum, fans, the room.

KA De-Noise reduces steady noise: hiss, hum, fans, the sound of a room. It compares the sound with a picture of the noise at every frequency and turns down what is only noise. Voice then shapes the tone of what is left and turns the voice up.

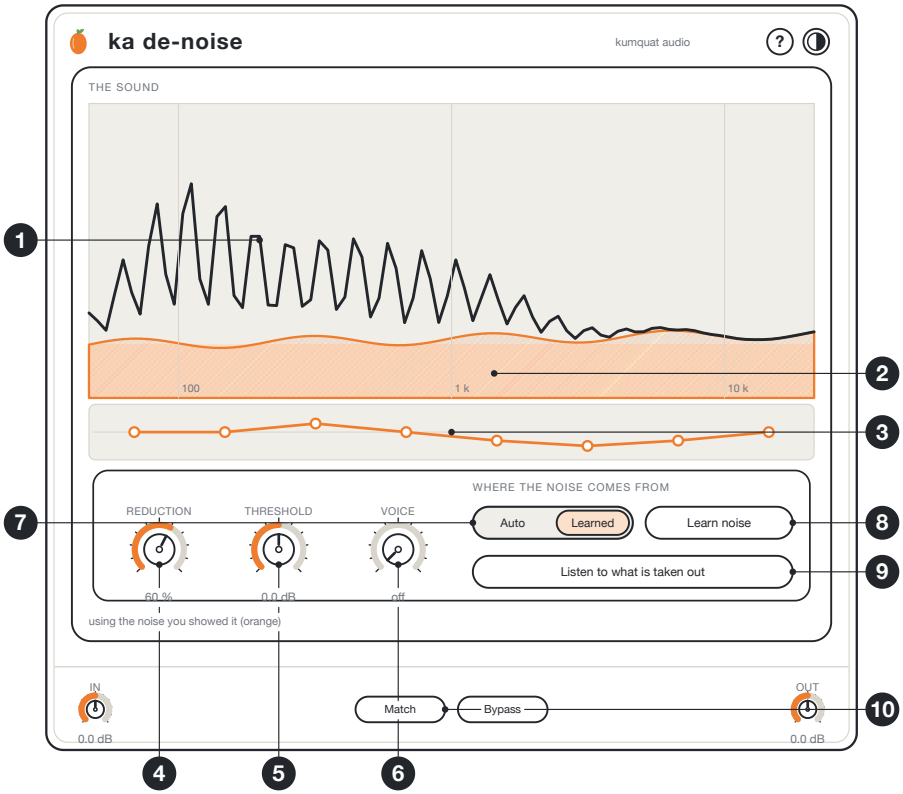


FIG. 7 The parts of KA De-Noise. The numbers are explained on the next page.

Parts

fig. 7

1. **The sound**
   A picture of the spectrum: the dark line is the sound, the orange area is the noise it has found.
2. **Noise estimate**
   Everything inside the orange area is treated as noise.
3. **Focus points**
   Pull a point up to take out more noise in that part of the spectrum, down to leave it more alone. Double-click a point to put it back.
4. **Reduction**
   How far the noise is turned down. Too much leaves a thin, watery voice.
5. **Threshold**
   Where the line between noise and sound is: up counts more as noise, down only what is plainly noise.
6. **Voice**
   Shapes the tone of the voice after the noise is out and turns the voice up where it is clearly voice. Off by default.
7. **Auto / Learned**
   Auto follows the noise by itself. Learned uses the noise you showed it.
8. **Learn noise**
   Press it while a noise-only stretch plays, wait two seconds or more, press it again. It stops by itself when the music stops, and then Learned is used.
9. **Listen**
   Plays only what is being taken out. It should be noise and nothing else.
10. **In, Out and Match**
    In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

Operation

step by step

1. Put the plug-in on the track. In Auto it finds the noise by itself in about a second.
2. Turn Reduction up until the noise has gone far enough. The picture shows the sound (dark line) and the noise it has found (orange).
3. Press Listen. You should hear only noise. If you hear the voice going thin or watery, lower Reduction, or lower Threshold so that less counts as noise.
4. Turn Voice up to improve the voice itself. Start at about 50 %. Listen does not play what Voice does, only what is taken out.
5. To take out more in one place than another, drag the points under the picture: pull one up where there is more noise to take out (hiss up high, rumble down low), pull one down to leave that part of the voice alone. The orange line follows.
6. For a recording with a long steady tone in it, or when you want it exact: play a stretch where nobody speaks, press Learn noise, wait at least two seconds, press it again. The plug-in switches to Learned by itself and uses what it heard.

Controls and settings

range · start

Control	Range	Starts at	What it does
Reduction	0 to 100 %	60 %	How far the noise is turned down.
Threshold	−12 to +12 dB	0 dB	Where the line between noise and sound is. Up: more counts as noise, so more is taken out. Down: only what is plainly noise.
Focus points	−12 to +12 dB each	0 dB	Eight points under the picture, from low to high frequencies: the same as Threshold, but for one part of the spectrum. Double-click a point to put it back.
Noise	Auto / Learned	Auto	Where the picture of the noise comes from. Auto follows the noise; Learned uses the noise you showed it.
Learn noise	button		Teaches it the noise: press, wait two seconds or more, press again. It stops by itself when the music stops (or after 20 seconds).
Voice	0 to 100 %	off	Moves the voice's tone toward well-recorded speech and turns it up where it is clearly voice (up to 6 dB at 100 %). Pauses are left alone.
Listen	on / off	off	Plays only what is being taken out.
In / Out	In ±24, Out ±12 dB	0 dB	The two knobs under the meters: before and after the plug-in.
Match	on / off	off	Keeps the output as loud as the sound that came in, whatever In and the plug-in do. See In, Out and Gain Match.

CAUTION Too much reduction leaves a thin, watery voice. Auto can be fooled by a long steady tone: use Learn.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo
Delay	2048 samples (43 ms at 48 kHz)
Window	760 × 620, can be resized

How it works

a closer look

KA De-Noise cuts the sound into about a thousand narrow slices of frequency, about two hundred times a second. For each slice it knows how loud the noise is, and it turns the slice down by as much as the noise accounts for. Where the voice is much louder than the noise, almost nothing is taken; where there is only noise, a lot is.

Auto keeps following the noise as the recording goes on. It needs about a second to find it, and a long steady tone can fool it. Learned uses the noise you showed it with Learn: it works at once, and it is exact.

Reduction sets how far down the noise is turned. Too far, and the parts of the voice that are as quiet as the noise are turned down with it, which is what makes a voice thin and watery.

A small neural network says, for each frequency and moment, how much of the sound is voice. While nobody speaks, the picture of the noise follows the sound freely; while someone speaks, it keeps away from the voice.

Voice works on what is left after the noise has gone. It compares the average tone of the voice over the last few seconds with that of well-recorded speech and adds the difference, up to 6 dB, only where the noise reduction kept the sound. It also turns the voice up where it is clearly voice, fast enough to be there at the start of a word.

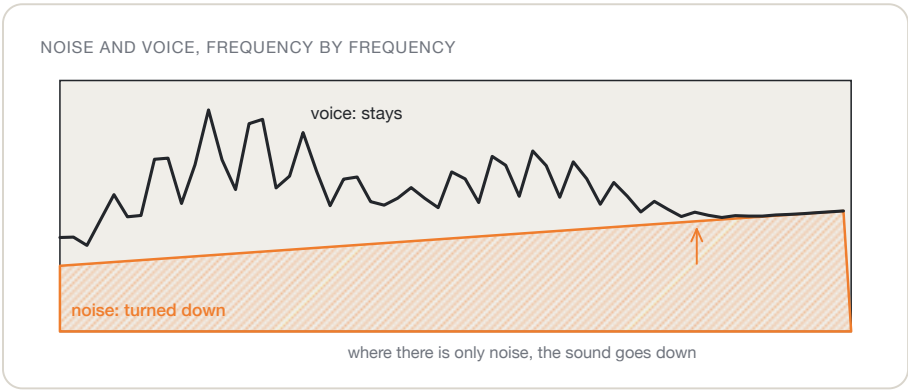


FIG. 8 Where there is only noise, the sound is turned down.

Credits and notices

KA De-Noise tells voice from noise with a small neural network. It was trained with recordings of speech, singing and noise: the speech from LibriSpeech (V. Panayotov, G. Chen, D. Povey and S. Khudanpur, from LibriVox recordings), the singing from VocalSet (J. Wilkins, P. Seetharaman, A. Wahl and B. Pardo), and some of the noise from MUSAN (D. Snyder, G. Chen and D. Povey). LibriSpeech, VocalSet and MUSAN are licensed under Creative Commons Attribution 4.0 (CC BY 4.0). The recordings are not part of the plug-in.

KA De-Noise also contains RNNoise, a library for the same job, which it can use in place of its own network. RNNoise is covered by the following license.

Copyright (c) 2007-2017, 2024 Jean-Marc Valin. Copyright (c) 2023 Amazon. Copyright (c) 2017, Mozilla. Copyright (c) 2005-2017, Xiph.Org Foundation. Copyright (c) 2003-2004, Mark Borgerding.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of the Xiph.Org Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NO. 4

KA De-Verb

Pulls a voice out of its room.

KA De-Verb pulls a voice out of its room. It takes away the echo and the wash that follow each sound, so the voice sounds closer and leaner. 0 is the sound as recorded.

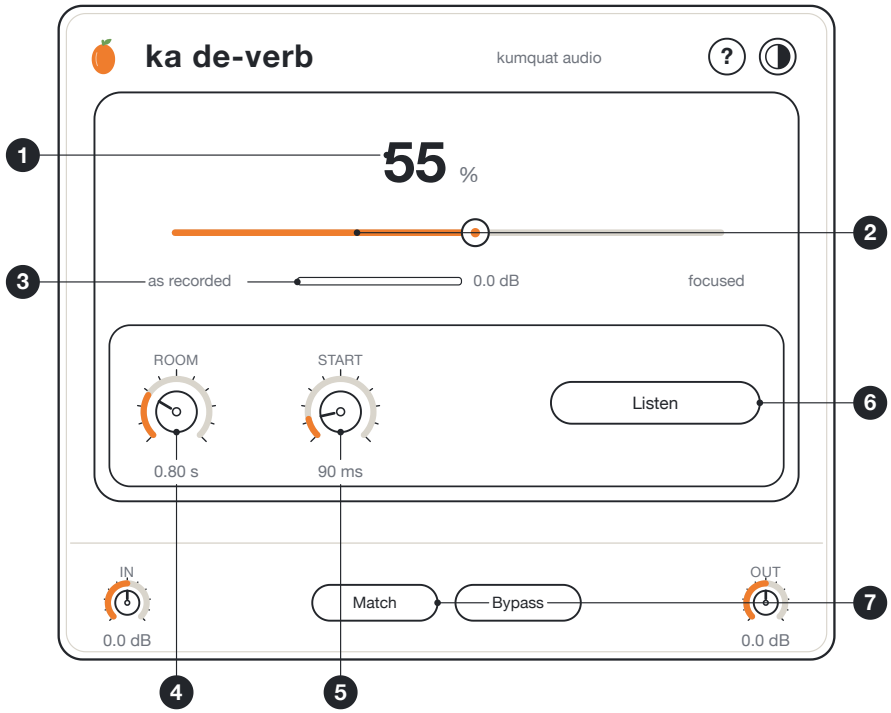


FIG. 9 The parts of KA De-Verb. The numbers are explained on the next page.

Parts

fig. 9

1. **Focus readout**
   How far the voice is being pulled out of its room.
2. **Focus**
   Left: as recorded. Right: the echo and wash behind the voice are taken away, so it sounds closer and leaner.
3. **Reduction bar**
   How much the sound is being turned down right now.
4. **Room**
   How long the room rings. Small rooms are short, halls are long.
5. **Start**
   How long after a sound begins the plug-in starts treating what follows as room.
6. **Listen**
   Plays only what is being removed.
7. **In, Out and Match**
   In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

Operation

step by step

1. Put the plug-in on the voice track and play a passage where you can hear the room.
2. Set Room to about how long the room rings: a small room is a few tenths of a second, a hall two or three.
3. Raise Focus until the room has gone far enough. The bar under the slider shows how much is being turned down.
4. Move Start to taste: later keeps more of the original sound and the first reflections.
5. Press Listen to hear only what is being removed.

Controls and settings

range · start

Control	Range	Starts at	What it does
Focus	0 to 100 %	50 %	How far the voice is pulled out of its room.
Room	0.2 to 3.0 s	0.7 s	How long the room rings.
Start	20 ms and up	130 ms	How long after a sound begins the plug-in starts treating what follows as room.
Listen	on / off	off	Plays only what is being removed.
In / Out	In ±24, Out ±12 dB	0 dB	The two knobs under the meters: before and after the plug-in.
Match	on / off	off	Keeps the output as loud as the sound that came in, whatever In and the plug-in do. See In, Out and Gain Match.

TIP If the voice sounds dry and dead, bring Focus down a little and Start up. A little room is part of a natural voice.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo
Delay	2048 samples (43 ms at 48 kHz)
Window	700 × 384, can be resized

How it works

a closer look

Every sound in a room is followed by its echo. KA De-Verb listens to how each frequency fades away after a sound, and takes away the part that is only ringing on, leaving the direct sound. That is why the voice comes closer.

Room tells it how long the ringing lasts, so that it knows what to look for. Start says how soon after a sound it may begin to take things away: the first reflections come just after the sound and are part of what makes it natural, so a later Start keeps more of them.

Focus is how much of the ringing is taken away. The bar under the slider shows it working. The plug-in looks ahead by 2048 samples to do this, and tells your music software.

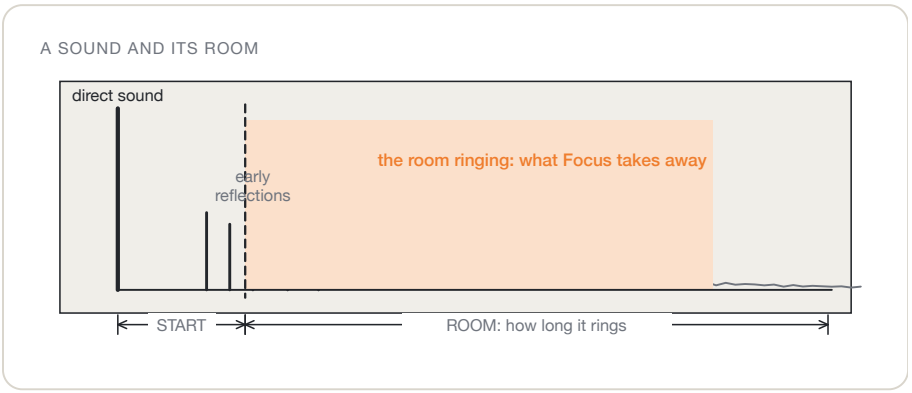


FIG. 10 A sound and its room: Start and Room set what is taken away.

NO. 5

KA EQ

Six bands, a resonance suppressor, and room made for a sidechain.

KA EQ has six bands and a high-pass and low-pass, and each band can be made dynamic so that it moves with the loudness around it. A resonance suppressor finds and turns down the bumps in a sound, and can make room for a track in the sidechain. The EQ, the suppressor and its sensitivity strip are all on one screen. The graph is the whole instrument: you work on it directly.

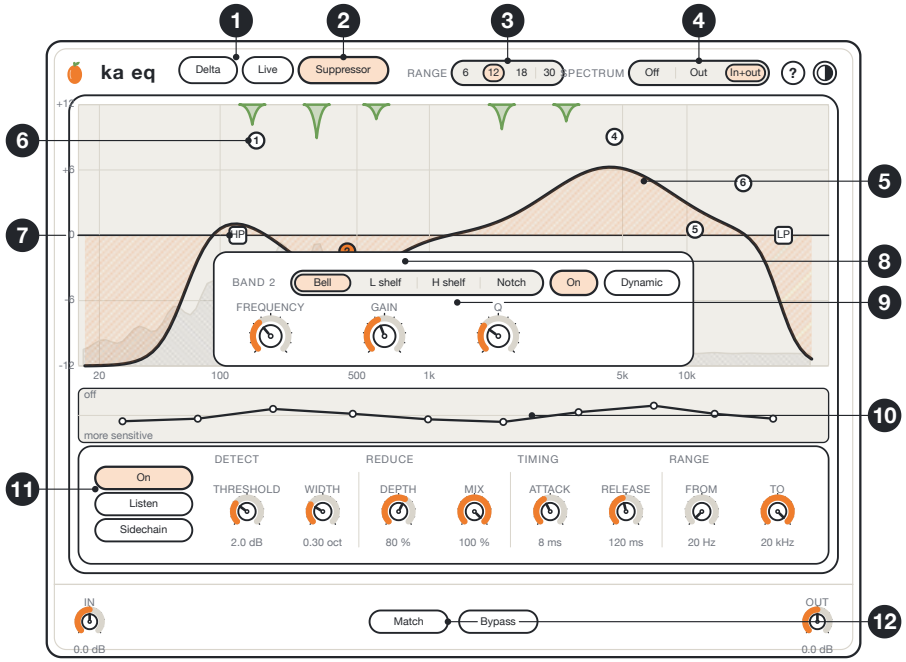


FIG. 11 The parts of KA EQ. The numbers are explained on the next page.

Parts

fig. 11

1. **Delta and Live**
   Delta plays only what the EQ is changing. Live switches the suppressor's look-ahead off, so the plug-in has no delay at all.
2. **Suppressor**
   Opens and closes the suppressor's controls under the graph. The On switch inside it is what turns the suppressor on.
3. **Range**
   How many dB the graph shows above and below the middle: 6, 12, 18 or 30.
4. **Spectrum**
   The live spectrum behind the curve: off, the output, or the input and the output.
5. **The graph**
   The curve of the EQ, over the spectrum. Click and drag on empty space to add a band. The green curtain hanging from the top is what the suppressor is taking out.
6. **A band (1 to 6)**
   Drag it for frequency and gain, scroll over it for Q. Resting the pointer on it shows its options.
7. **HP and LP**
   The high-pass and low-pass: they show on the graph only when they are on. Press one to switch it on, drag it to move it, scroll over it for its slope.
8. **Option bar**
   Appears beside a band under the pointer: type, Dynamic, and × to remove it.
9. **The selected band**
   A small card at the bottom of the graph: its type, On, Dynamic (which makes the band move with the loudness around it), and its frequency, gain and Q (and, when it is dynamic, its range, threshold, attack and release).
10. **Sensitivity strip**
    Where the suppressor is touchy, from the sub-bass (30 Hz) to the air (12 kHz), on the same scale as the graph. Drag a point DOWN to make that part of the spectrum act more readily, UP to leave it alone more.
11. **The suppressor**
    On switches the resonance suppressor on (it adds a little delay, which the plug-in tells your music software). Listen plays only what it is taking out. Sidechain turns this sound down where the track in the sidechain has energy. The knobs are in four groups: Threshold and Width (what counts as a resonance), Depth and Mix (how much is taken out), Attack and Release (how fast), From and To (where it works).
12. **In, Out and Match**
    In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

The EQ

working on the graph

1. To add a band, click on empty space in the graph and drag. The new band starts where you pressed and follows the pointer. Hold Option while you do it to make a dynamic band.
2. Drag a band for its frequency and gain. Scroll (two fingers on a trackpad, the wheel on a mouse) over it, or with the band selected, to change its Q: up narrows it, down widens it. On HP and LP, scrolling sets the slope.
3. Rest the pointer on a band: a small bar appears beside it with the band's type, Dynamic, and × to remove it.
4. Hold Control and press on a band, the high-pass or the low-pass to hear only what it works on (for the high-pass and low-pass, what they take away); the same on a point of the suppressor's sensitivity strip to hear its part of the spectrum. Double-click a band to delete it. Option-click a band to make it dynamic, or not; hold Option while dragging, or while scrolling, to set how much dynamic EQ it has (scrolling with Option makes the band dynamic if it is not).
5. Press HP or LP (in the bottom corners of the graph) to switch it on. Once on, it shows on the graph: drag it to move it, scroll over it to set the slope.
6. Right-click a band for its menu.
7. The band's own controls float over the graph beside the band when you select it. The suppressor's controls open under the graph from the Suppressor button at the top: switch it on with On (see the suppressor pages).

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo, with a sidechain input
Delay	none; 4096 samples (85 ms) with the suppressor on, and none with Live on
Window	1000 × 578, can be resized

Controls and settings

range · start

Control	Range	Starts at	What it does
Band 1–6	20 Hz – 20 kHz	off	Bell, low shelf, high shelf or notch; frequency, gain ±24 dB, Q.
High-pass	20 Hz – 1 kHz	80 Hz, off	Slopes of 6, 12, 18, 24, 36, 48, 72 or 96 dB per octave.
Low-pass	1 kHz – 20 kHz	12 kHz, off	The same slopes.
Dynamic range	–24 to +24 dB	–6 dB	How far a dynamic band moves when the sound around it is loud.
Dynamic threshold	–60 to 0 dB	–30 dB	How loud it must get before the band moves.
Dynamic attack / release	0.5–100 ms / 10–1000 ms	8 / 120 ms	How fast the band moves, and lets go.
Range	6, 12, 18, 30 dB	12	What the graph shows above and below the middle.
In / Out	In ±24, Out ±18 dB	0 dB	The two knobs under the meters: before and after the plug-in.
Match	on / off	off	Keeps the output as loud as the input. See In, Out and Gain Match.
Live	on / off	off	Switches the suppressor off, so there is no delay at all.

TIP For playing through the EQ, switch Live on: no delay, and the suppressor is off until you switch Live off.

TIP A dynamic band with a negative range and a narrow Q is a de-esser; with a wide Q it is a gentle way of keeping a boomy voice under control.

How it works

the EQ

Each of the six bands, and the high-pass and low-pass, is a smooth filter that adds no delay. A band can be a bell, a shelf or a notch. Dragging a band on the graph changes its frequency and gain; scrolling changes its Q.

A dynamic band has a second detector that listens to the sound around the band. When that sound gets louder than Threshold, the band's gain moves, by up to Range: a negative range makes a band that cuts when it is loud (a de-esser), a positive one a band that opens up.

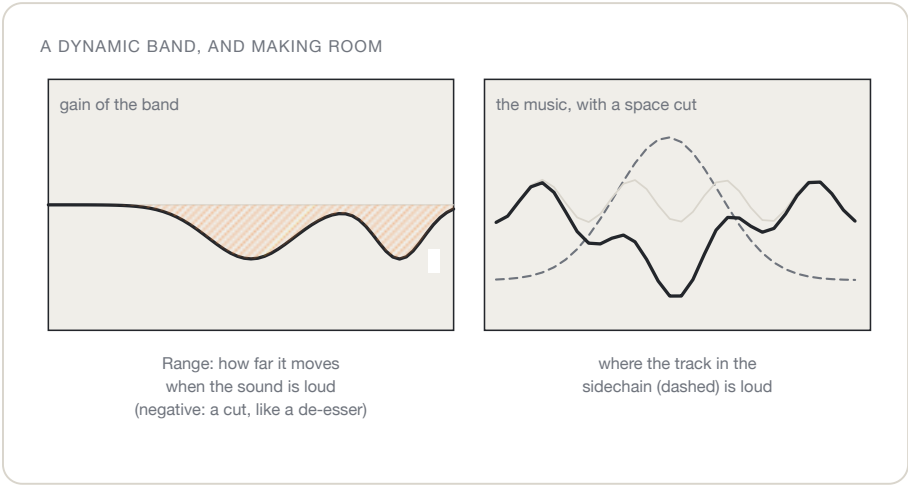


FIG. 13 A dynamic band moves with the sound; The sidechain cuts a space for another track.

The resonance suppressor

and making room

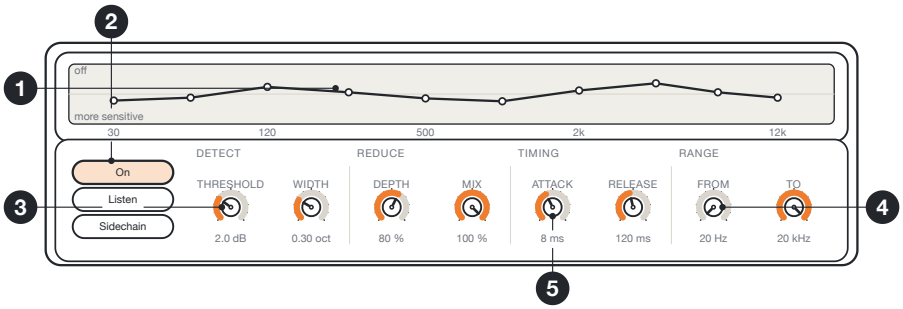


FIG. 12 The suppressor view: what it takes out hangs from the top of the graph, in green.

1 Sensitivity strip
Ten points from 30 Hz to 12 kHz. Drag a point DOWN to make that part of the spectrum more sensitive (the suppressor acts more readily), UP for less; at the very top it leaves that part alone. Dragging across the strip paints. Double-click a point to reset it.

2 On, Listen, Sidechain
On switches the suppressor on. Listen plays only what it takes out; Sidechain makes room for the track in the sidechain.

3 Threshold, Width, Depth, Mix
How far a bump must stick out before it is touched, how wide a bump may be (down to 0.01 octave), how much of each resonance is taken out, and how much of the processed sound is mixed with the original.

4 From and To
The range it works in. They can cover the whole spectrum: From goes down to 20 Hz, To up to 20 kHz.

5 Attack and Release
How fast the reduction comes in, and how fast it lets go (down to 5 ms).

Using the suppressor

step by step

1. Press Suppressor at the top to open its controls under the graph, then On to switch it on.
2. Raise Depth until the resonances are tamed. Lower Threshold to find more of them; make Width smaller to be more exact.
3. Press Listen (in the suppressor's controls) to hear only what is being taken out, and check that it is the right thing.
4. Use From and To to say where it works (they can cover the whole spectrum, from 20 Hz to 20 kHz). Drag the points of the sensitivity strip under the graph to make a part of the spectrum act more readily (down) or less (up), all the way up to leave it alone. The first two points are the sub-bass, at 30 and 60 Hz.
5. To make room for another sound (music under dialogue): send the dialogue track to this plug-in's sidechain in Logic, and press Sidechain. The music is turned down where the dialogue has its energy.

Control	Range	Starts at	What it does
Depth	0 to 100 %	30 %	How much of each resonance is taken out.
Threshold	0 to 24 dB	6 dB	How far a bump must stick out before it is touched.
Width	0.01 to 1 oct	0.30 oct	How wide a bump may be.
Attack	1–50 ms	8 ms	How fast the reduction comes in.
Release	5–500 ms	120 ms	How fast it lets go.
From / To	20–2000 Hz / 2–20 kHz	20 Hz / 20 kHz	The range it works in: the whole spectrum, to begin with.
Sensitivity strip	10 points, 30 Hz to 12 kHz	flat	Down: that part of the spectrum is treated more readily. Up: less, and at the very top not at all.
Mix	0 to 100 %	100 %	How much of the processed sound is mixed with the original.

How it works

the suppressor

The suppressor looks at the sound in slices of 4096 samples (85 ms) and compares each narrow bump with its wider surroundings. A bump that sticks out by more than Threshold is turned down by Depth, but only between HP and LP. Make room does the same with a second sound: the spectrum of the sidechain says where to cut.

NO. 6

KA Limiter

A true-peak limiter with a clipper, three modes and a little bite.

KA Limiter holds a mix at a ceiling while making it as loud as you ask, without squashing it. A clipper in front shaves the highest peaks, three modes change how the limiter works inside, and a small control called Bite lifts the mids for presence. A picture of what the clipper and the limiter take off shows you what is happening as it happens.

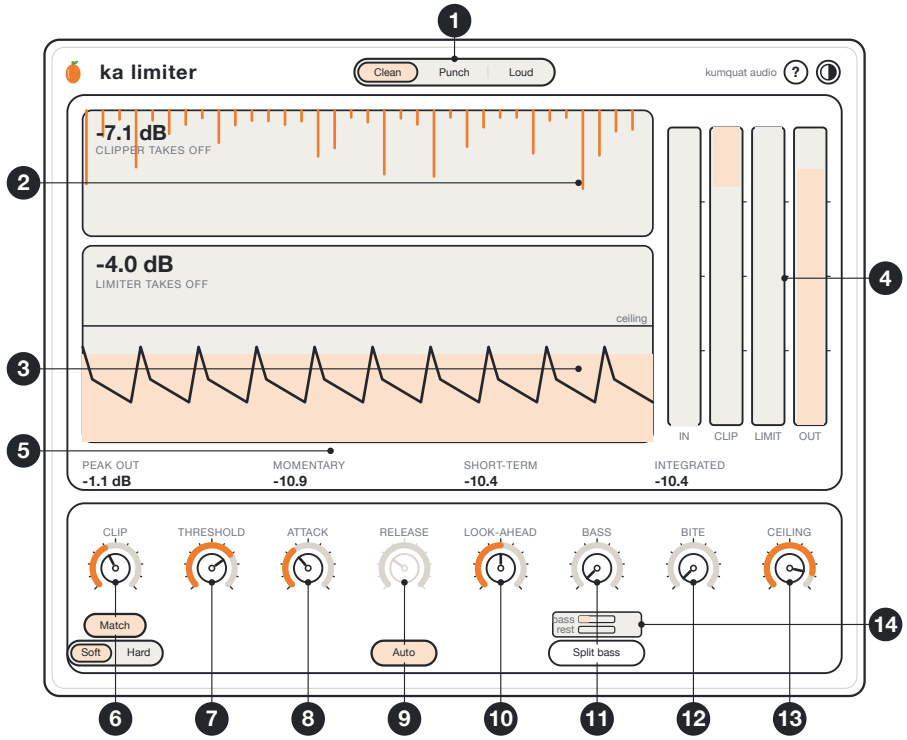


FIG. 14 The parts of KA Limiter. The numbers are explained on the next page.

Parts

fig. 14

1. **Mode**
   Clean, Punch or Loud: three ways of working inside. The knobs stay put.
2. **Clipper picture**
   What the clipper takes off the peaks, over the last few seconds.
3. **Limiter picture**
   What goes in, what the limiter takes off (grey), what comes out (orange).
4. **Meters**
   IN, CLIP, LIMIT and OUT, with a held-peak line. Click to clear.
5. **Loudness**
   Peak out, momentary, short-term, integrated. Click integrated to restart.
6. **Clip**
   Drives the sound into a clipper that shaves the highest peaks. Soft or Hard; Match.
7. **Threshold**
   Turn it down to push the sound into the limiter.
8. **Attack**
   How fast it reacts to sound that stays loud.
9. **Release**
   How fast it lets go (with Auto it chooses).
10. **Look-ahead**
    How far ahead the limiter looks.
11. **Bass**
    With Split bass on: how much low end is kept in.
12. **Bite**
    A smooth, very wide lift of the mids and high mids.
13. **Ceiling**
    The highest peak that comes out.
14. **Bass picture**
    How much the bass and the rest are being turned down.

Operation

step by step

1. Put the plug-in last on the master. Set Ceiling (–1 dB is a good place to begin).
2. Turn Threshold down to push the sound into the limiter. It gets louder and the limiter works harder.
3. Add a little Clip (a few dB, with Match on) if the peaks are doing the harm. The upper picture shows what it takes off.
4. Choose a mode: Clean for the most transparent, Punch to keep the snap of the hits, Loud for the most volume.
5. Leave Release on Auto. Shorten Look-ahead for snappier hits; lengthen it to be smoother, most of all on the bass.
6. If the kick pulls the mix down, turn Split bass on and raise Bass. The two small bars show who takes how much.
7. Raise Bite a little for presence, and compare with it off: the ceiling is the same.
8. Press Compact (top right) to shrink the window to two meters, the Clip and Threshold knobs and the loudness readings, and leave it open on the master while you mix. Press Full to come back.

Controls and settings

range · start

Control	Range	Starts at	What it does
Mode	Clean / Punch / Loud	Clean	How it works inside: three engines. The knobs are left alone.
Clip	0 to 20 dB	off	Drives the sound into a clipper that shaves the highest peaks. Soft or Hard; Match brings the level back down.
Threshold	–20 to 0 dB	0 dB	How far the sound is pushed into the limiter.
Attack	0.05 ms to 10 s	15 ms	How fast it reacts to sound that stays loud. At seconds, only the brick wall works.
Release	1 ms to 10 s	150 ms	How fast it lets go. Works when Auto is off.
Look-ahead	0.5 to 5 ms	2.5 ms	How far ahead it looks.
Bass	0 to 100 %	0 %	With Split bass: how much low end is kept in.
Bite	0 to 100 %	off	A smooth, very wide lift of the mids and high mids, centred on 3 kHz.
Ceiling	–6 to 0 dB	–1 dB	The highest peak that comes out (true peak).
Compact / Full	button	Full	Compact is a small window with the Clip and Limit meters, the Clip and Threshold knobs and the loudness readings (momentary, short-term, integrated). Full is the whole window.

CAUTION Release of a millisecond or two, with a long Attack, is the dirtiest setting there is: the gain moves inside every bass wave. Use Auto, or 60 to 100 ms.

TIP In Punch and Loud, part of the reduction is done by a soft clipper: the readouts count it, so they show the whole of the limiting.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo
Delay	746 samples (15.5 ms at 48 kHz), whatever Look-ahead is
Ceiling	true peak, with an exact last limiter
Window	860 × 688, can be resized[34]

How it works

the signal path

The sound goes through Bite (if it is on), then Clip, then Threshold, which turns it up into the limiter, then the limiter, then the ceiling. The ceiling is a true peak: the height of the wave between the samples, which is what a converter or an encoder makes of it.

Inside the limiter there are two steps. A slow one reacts only to sound that stays loud, so a drum hit hardly moves it. Then a look-ahead brick wall takes what is left over the ceiling: it finds the lowest gain needed in the next few milliseconds, lets go with a release that is longer the longer it has been working, and smooths the result so that the gain never has a corner.

In Clean mode that is all there is: only the gain moves. In Punch and Loud the limiter is allowed to leave the very tips of the peaks a little over the ceiling, and a soft clipper takes them off, instead of the gain ducking everything around them. A short last limiter makes sure the ceiling still holds exactly.

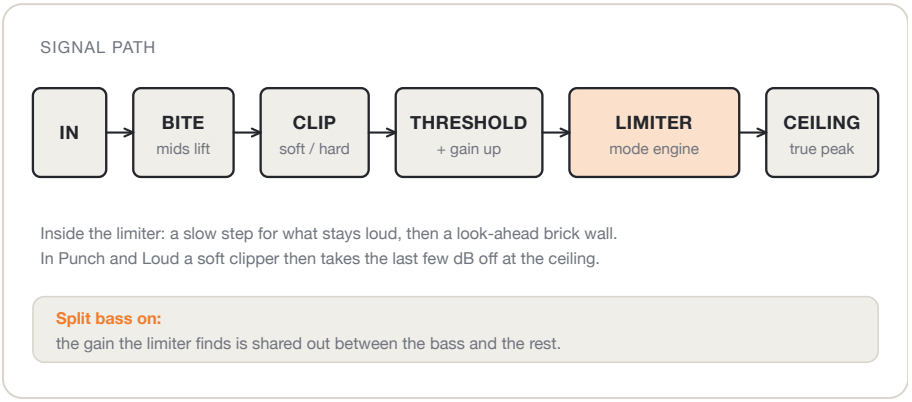


FIG. 15 The signal path.

The modes

and the bass

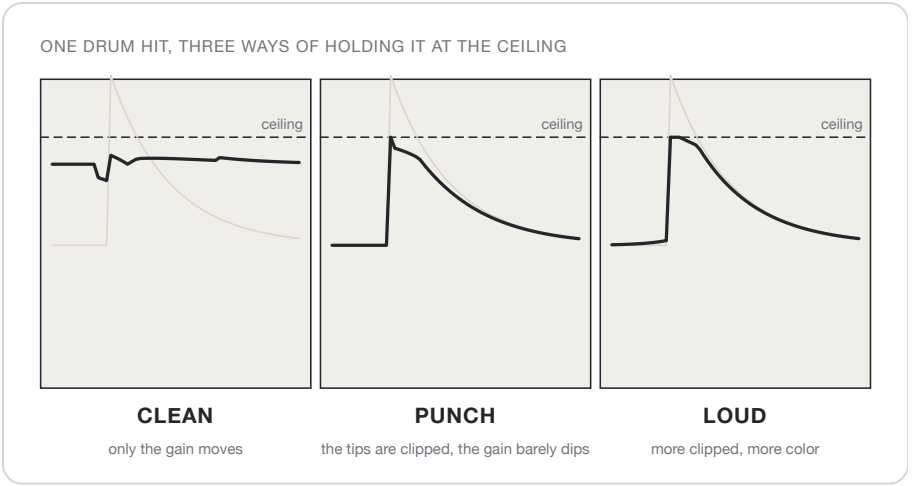


FIG. 16 The three modes hold the same hit at the same ceiling in different ways.

Split bass divides the sound at about 140 Hz. The limiter still works out one gain from the whole sound, but the bass is given less of it, and the rest takes the difference, worked out from how much of the peak each part is. So the sound still stops exactly at the ceiling, while the kick no longer pulls the voice down. Bass at 0 % is the same as having Split bass off.

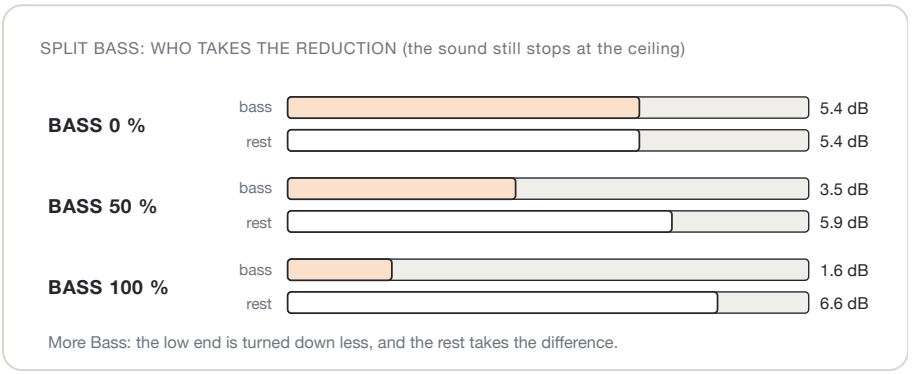


FIG. 18 Split bass moves the reduction onto the bass; the ceiling still holds.

Bite

presence without level

Bite is four wide, smooth bells: a very wide one at 3 kHz (Q 0.2, about five octaves across), a small one at 1.6 kHz, one at 6.5 kHz, and a small scoop in the low mids. It lifts the range where the ear is most sensitive, so a mix feels louder and more present at the same ceiling. It adds no delay.

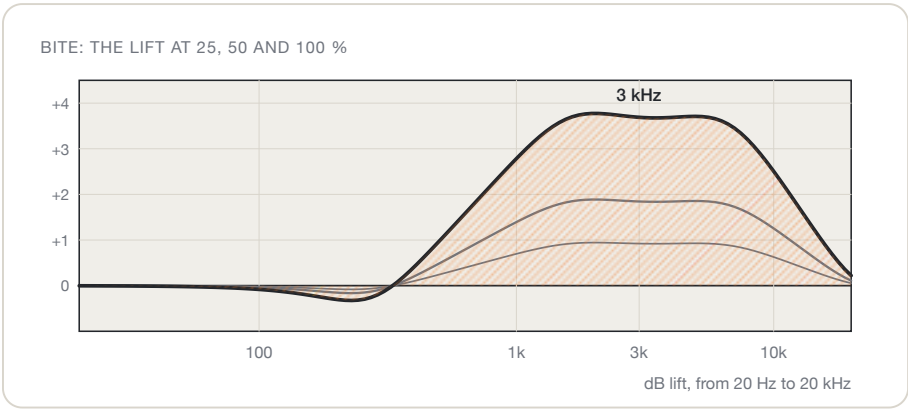


FIG. 17 Bite: a wide, smooth lift, and the bass is left alone.

Places to start

then adjust by ear

For	Mode	Try	You get
A gentle master	Clean	Threshold –2 to –4 dB; Clip off; Release Auto; Bite off	A few dB of reduction. The most transparent.
Loud, with snap	Punch	Threshold –4 to –6 dB; Clip 2 to 4 dB with Match on; Look-ahead 1.5 to 2.5 ms	Louder, and the hits keep their front.
As loud as it goes	Loud	Threshold –6 dB or more; Clip 4 dB, Hard if you like the bite	The most volume, and the most color.
A kick that pulls the mix down	any	Split bass on; Bass 30 to 60 %	The low end is turned down less, the rest more.
A mix that needs presence	any	Bite 20 to 50 %	Compare with it off: the ceiling is the same.

NO. 7

KA Mic Blend

Blends two mics on one source, and lines them up.

KA Mic Blend mixes two microphones on one source with a single control, and lines them up. When two mics hear the same sound at different distances, one arrives late, and if it is also upside down the two cancel each other and the sound goes thin. Align measures this and sets it right.



FIG. 19 The parts of KA Mic Blend. The numbers are explained on the next page.

Parts

fig. 19

1. **A and B**
   How much of each mic is in the blend. With a stereo track A is the left channel and B the right; with a sidechain, A is the track and B the sidechain.
2. **Sway**
   The blend: all the way left is only A, all the way right only B.
3. **B offset**
   Delays B so that it lines up with A. Align sets it.
4. **Flip B polarity**
   Turns B upside down. Align sets it when the two mics fight each other.
5. **Blend law**
   Linear for mics that are in time, Equal power for mics that are apart.
6. **Align**
   Play the source, then press it: it measures the timing and polarity of A and B.
7. **Alignment on**
   Switches the alignment off and on so you can hear what it does.
8. **Status**
   What the alignment is doing, or what it found.
9. **Phase meter**
   Shows whether the two mics add up (in phase) or cancel (out of phase).
10. **In and Out**
    In turns both mics up or down before the blend. Out turns the blend up or down.

Operation

step by step

1. Put the plug-in on a stereo track that holds the two mics (A is the left channel, B the right), or on a track with the other mic in the sidechain (A is the track, B the sidechain).
2. Play the source and press Align. It listens for a moment, analyses, and sets B offset and Flip B polarity.
3. Move the Sway slider to blend the mics: all the way left is only A, all the way right only B.
4. Switch Alignment off and on to hear what it does. What was measured is kept.
5. Choose Linear for mics that are in time, Equal power for mics that are apart.

Controls and settings

range · start

Control	Range	Starts at	What it does
Sway	–100 to +100 %	0	The blend of A and B.
B offset	–10 to +10 ms	0	Delays B to line it up with A. Align sets it.
Flip B polarity	on / off	off	Turns B upside down. Align sets it.
Alignment on	on / off	on	Switches B offset and the flip off and on.
Blend law	Linear / Equal power	Linear	How the two add up as you move the slider.
In / Out	In ±24, Out –24 to +12 dB	0 dB	In turns both mics up or down before the blend; Out turns the blend.

TIP Play something with a clear start while it listens: a drum, a word, a strummed chord.

Specifications

Formats	Audio Unit, VST3
Channels	stereo, or one track and a sidechain
Delay	484 samples (10 ms at 48 kHz): fixed, so that Align can move B either way
Window	700 × 420, can be resized

How it works

a closer look

When two mics hear the same sound at different distances, the sound reaches one later than the other. Blended together, the two copies comb against each other and the sound goes thin; if one mic is also upside down to the other, they cancel.

Align plays the source and compares A and B at every offset up to ten milliseconds, either way, and in both polarities. It picks the setting where the two agree most, sets B offset and Flip B polarity, and keeps what it found.

Because B might have to move earlier or later than A, the plug-in always delays both by the longest offset (about 10 ms) and moves B within that, so the delay your music software is told about never changes.

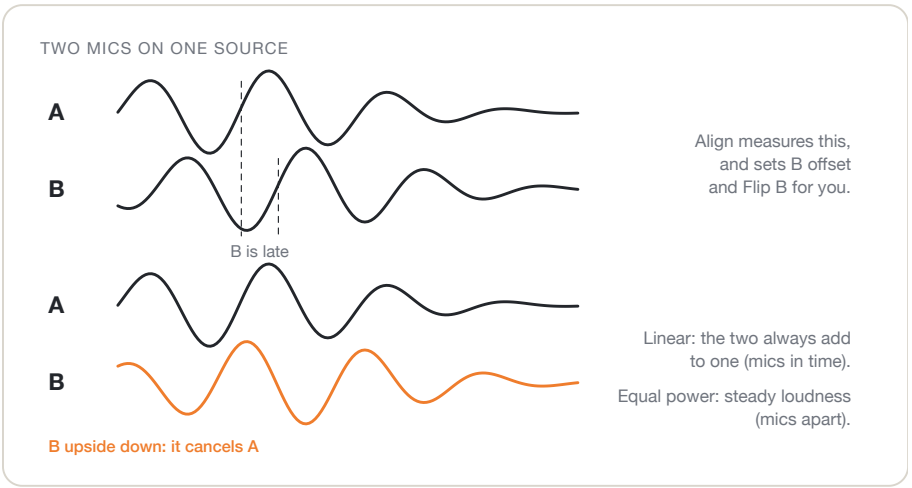


FIG. 20 What Align measures: the delay of B, and whether it is upside down.

NO. 8

KA Drum Gate

Finds every hit on every drum mic, and gates and cleans each drum.

KA Drum Gate reads a whole drum recording, finds every hit on every mic, and then gates and cleans each drum. A snare is also heard, quieter, in the kick mic, so a hit is counted only in the mic it belongs to. It works on whole recordings, so it reads the audio as it plays.

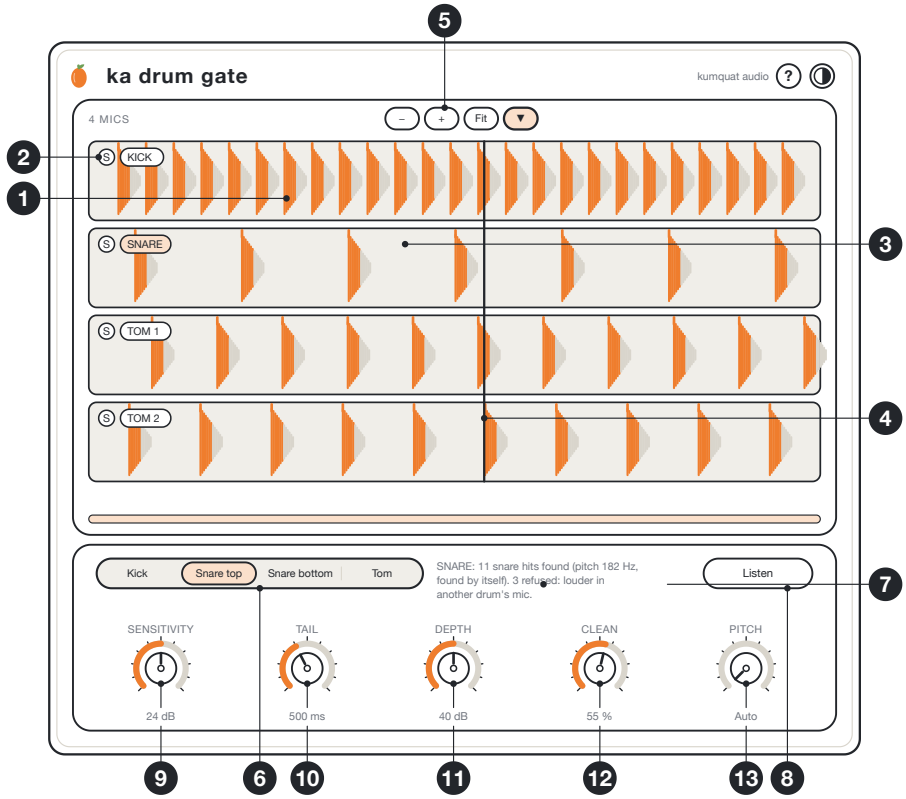


FIG. 21 The parts of KA Drum Gate. The numbers are explained on the next page.

Parts

fig. 21

1. **A lane**
   One for each drum mic: faint strokes are what it hears, orange the hits found in it.
2. **S and name**
   S solos the drum. Click a name to pick that drum and work on it.
3. **A hit**
   The orange tick over each hit marks where it was found.
4. **Playhead**
   Click to move it; scroll to move along, pinch or Cmd+scroll to zoom.
5. **Zoom**
   Zoom out, zoom in, fit, and follow the playhead.
6. **Drum type**
   The kind of drum this mic is (guessed from the track name).
7. **Message**
   What was found, and what was refused (louder in another drum's mic).
8. **Listen**
   Plays only what is being taken out.
9. **Sensitivity**
   How quiet a hit may be and still count. Higher finds ghost notes.
10. **Tail**
    The longest the drum may ring before the gate closes.
11. **Depth**
    How far the sound is turned down between hits.
12. **Clean**
    Takes cymbal and bleed out of the drum while its gate is open.
13. **Pitch**
    The drum's pitch. Auto finds it from the recording.

Operation

step by step

1. Open Logic Pro using Rosetta (see the page headed “Plug-ins that read whole recordings”).
2. Put the plug-in first in the Audio FX list on each drum mic’s track. Select a region and press play.
3. Wait while it reads. The picture fills with one lane for each mic, and an orange tick over each hit.
4. Click a drum’s name to work on it. Check the drum type is right (it is guessed from the track name).
5. Raise Sensitivity to find quieter hits (ghost notes). Set Tail to how long the drum may ring, and Depth to how far the sound is turned down between hits.
6. Raise Clean to take cymbal and bleed out of the drum while its gate is open. Press S to hear one drum alone, and Listen to hear what is being taken out.

Controls and settings

range · start

Control	Range	Starts at	What it does
Sensitivity	6 to 48 dB	24 dB	How quiet a hit may be and still count.
Tail	20 ms to 3 s	500 ms	The longest the drum may ring before the gate closes.
Depth	0 to 80 dB	40 dB	How far the sound is turned down between hits.
Clean	0 to 100 %	55 %	Takes cymbal and bleed out of the drum while its gate is open.
Pitch	Auto, 30 to 2000 Hz	Auto	The drum's pitch.
Drum type	Kick, Snare top, Snare bottom, Tom	guessed	What kind of drum this mic is.
S / Listen	buttons		Solo one drum; hear what is being taken out.

TIP Put it on the kick and the snare tracks too, not only the toms. It uses every mic to tell whose hit is whose.

CAUTION Not on the overheads. They hear the whole kit at once and have no hits of their own.

Specifications

Formats	Audio Unit only (needs ARA)
Hosts	Logic Pro, run with Rosetta
Channels	mono or stereo
Window	800 × 650, can be resized

How it works

a closer look

KA Drum Gate reads each drum mic's whole recording as it plays (this is what ARA is for) and finds where the hits are. It uses the drum's pitch, from the track name or from the recording itself, and Sensitivity, which says how quiet a hit may be and still count.

Then it compares each hit with the other mics. A snare hit is heard in the kick mic too, but quieter: the hit belongs to the mic where it is loudest, and is refused in the others. This is why it works best with a plug-in on every mic.

For each drum there is then a gate that opens at every hit and closes after Tail, turning the sound down by Depth in between. While the gate is open, Clean takes cymbal and bleed out of the drum.

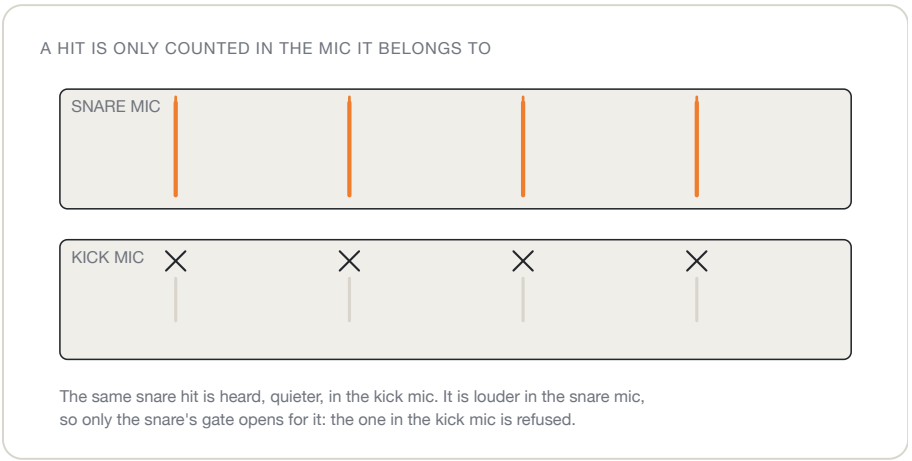


FIG. 22 Bleed: the same snare hit is heard in two mics, and counted in one.

NO. 9

KA Silence Trim

Finds the silences in a recording and keeps what matters.

KA Silence Trim finds the silences in a recording and keeps what matters. It reads the whole clip, shows what it would keep, and writes a trimmed file you can drag into your project. The file is as long as the clip, with the cut parts silent, so everything stays in time.

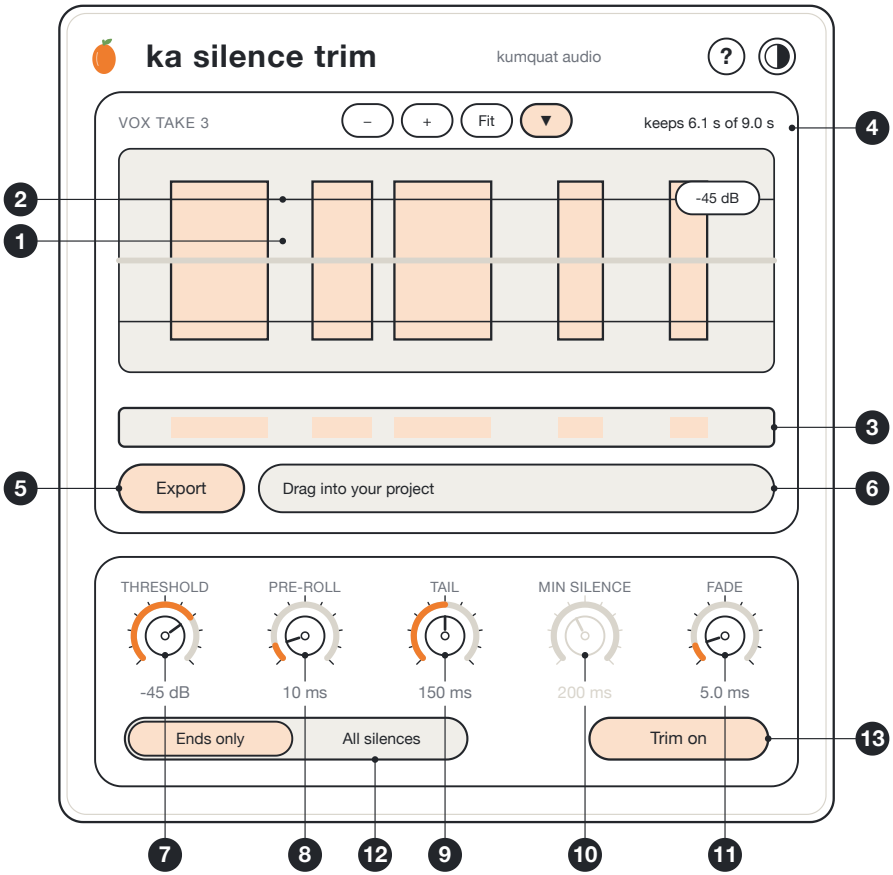


FIG. 23 The parts of KA Silence Trim. The numbers are explained on the next page.

Parts

fig. 23

1. **Wave view**
   The recording: orange is what is kept, the faint band is the silence that is cut.
2. **Threshold line**
   The thin lines mark the threshold. Drag them up and down.
3. **Overview**
   The whole clip at a glance; drag it to move along.
4. **Readout**
   How much is kept, of how much there is.
5. **Export**
   Writes the trimmed audio to a file now.
6. **Drag handle**
   Pick it up and drop it into your project: the file is written as you pick it up.
7. **Threshold**
   Anything quieter than this counts as silence.
8. **Pre-roll**
   How much to keep before a sound starts, so the attack is never clipped.
9. **Tail**
   How much to keep after a sound ends, so it can ring out.
10. **Min silence**
    With All silences: a gap shorter than this is left alone.
11. **Fade**
    The length of the fade at each cut.
12. **Ends only / All silences**
    Cut only the silence at the start and the end, or the gaps in the middle too.
13. **Trim on**
    Turn it off to hear the original recording.

Operation

step by step

1. Open Logic Pro using Rosetta, and put the plug-in first in the Audio FX list of an audio track.
2. Select a clip and press play. The wave view fills in, with orange where it keeps the sound.
3. Set Threshold (or drag the threshold lines in the picture): anything quieter than this counts as silence. The readout shows how much is kept.
4. Set Pre-roll and Tail so that attacks are never clipped and sounds can ring out. Fade sets the small fade at each cut.
5. Choose Ends only (cut the silence at the start and end) or All silences (cut the gaps in the middle as well; Min silence says how long a gap must be).
6. Press Export, or pick up the drag handle and drop it in your project. The file is written to Music/Snip.

Controls and settings

range · start

Control	Range	Starts at	What it does
Threshold	–70 to –10 dB	–45 dB	Anything quieter counts as silence.
Pre-roll	0 to 200 ms	10 ms	How much to keep before a sound starts.
Tail	0 to 1000 ms	150 ms	How much to keep after a sound ends.
Min silence	20 to 2000 ms	200 ms	With All silences: a shorter gap is left alone.
Fade	1 to 50 ms	5 ms	The length of the fade at each cut.
Mode	Ends only / All silences	Ends only	Which silences are cut.
Trim on	on / off	on	Off plays the original recording.

TIP Drop the exported file at the start of the clip. If you want slices, use Strip Silence in Logic on it afterwards.

Specifications

Formats	Audio Unit only (needs ARA)
Hosts	Logic Pro, run with Rosetta
Output	one WAV file, as long as the clip
Window	700 × 620, can be resized

How it works

a closer look

KA Silence Trim reads the whole clip and measures its level. Anything below Threshold counts as silence. For every stretch of sound it keeps Pre-roll before it and Tail after it, and puts a small Fade at each cut.

With Ends only, only the silence before the first sound and after the last one is cut. With All silences, the gaps in the middle are cut too, unless they are shorter than Min silence.

The exported file is as long as the clip, with the cut parts silent, and it is written to Music/Snip. Because nothing is moved, it lines up with the original: drop it at the start of the clip.

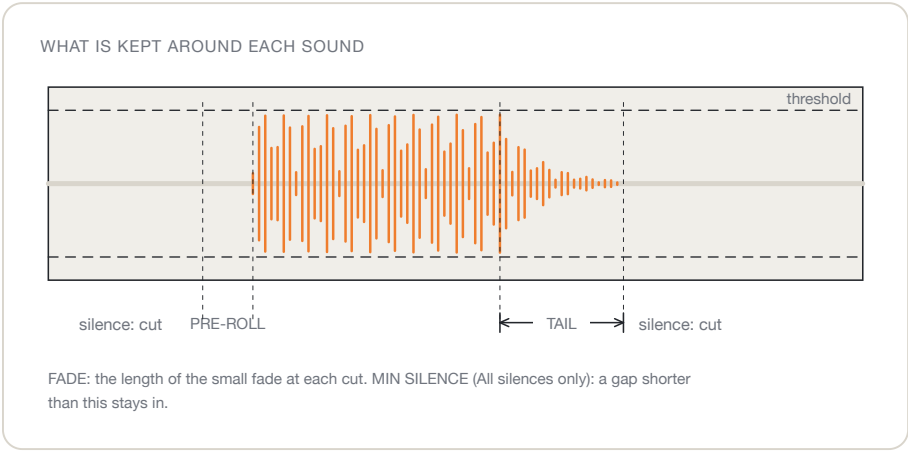


FIG. 24 What is kept around each sound.

NO. 10

KA Auto Mixer

Turns up whoever is speaking, on several mics.

KA Auto Mixer turns up whoever is speaking, on several mics, and turns the others down. It reads all the tracks, works out who is talking when, and rides the faders for you. It also checks which mic hears each sound first, so one person's voice in another person's mic does not open it, and a lead mic wins when it is a close call.

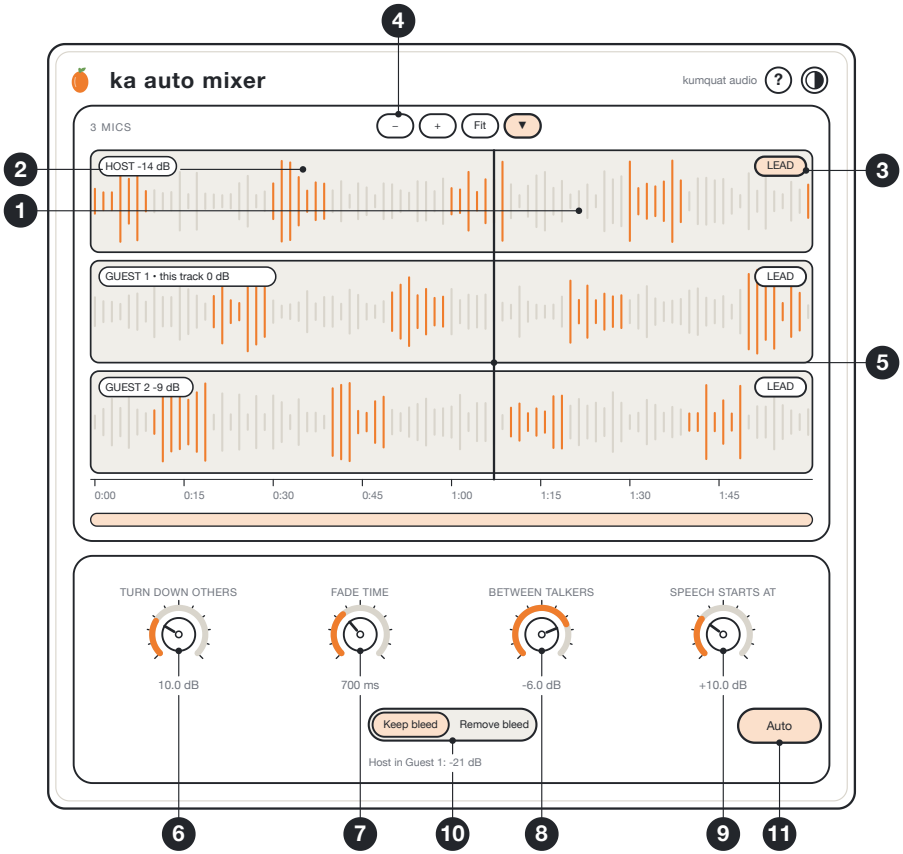


FIG. 25 The parts of KA Auto Mixer. The numbers are explained on the next page.

Parts

fig. 25

1. **A lane**
   One for each mic. Faint strokes are what the mic hears; orange is what KA Auto Mixer lets through.
2. **Name and level**
   The track's name and the dB it is being turned by now. • this track marks the one you are on.
3. **LEAD**
   Marks a lead mic (the host's, say): when it is a close call who is talking, the lead mic wins.
4. **Zoom**
   Zoom out, zoom in, fit, and follow the playhead.
5. **Playhead**
   Click anywhere to move it. Scroll to move along, pinch or Cmd+scroll to zoom.
6. **Turn down others**
   While someone speaks, the other mics go down by this much.
7. **Fade time**
   How long a mic takes to fade up or down when the person talking changes.
8. **Between talkers**
   How far the other mics go down while nobody speaks. The last speaker's mic stays open.
9. **Speech starts at**
   How far above its own noise a mic must be to count as someone talking.
10. **Keep / Remove bleed**
    Remove bleed (experimental) takes a speaker's voice out of the other mics. Under it: the bleed that was found.
11. **Auto**
    Listens to your recordings and sets everything for you.

Operation

step by step

1. Open Logic Pro using Rosetta, and put the plug-in first in the Audio FX list on every mic’s track.
2. Select the regions and press play. The picture fills with one lane for each mic. A track with no KA Auto Mixer on it is marked in red.
3. Press Auto: it listens to your recordings and sets Turn down others and Speech starts at for you.
4. Adjust Turn down others (how far the other mics go down while someone speaks) and Fade time (how long a mic takes to come up or go down).
5. Mark a lead mic with LEAD, the host’s for example: when it is a close call who is talking, the lead mic wins.
6. Between talkers sets what happens in the pauses: the mic of the last speaker stays open and the others go down by this much.

Controls and settings

range · start

Control	Range	Starts at	What it does
Turn down others	3 to 24 dB	10 dB	How far the other mics go down while someone speaks.
Fade time	100 ms to 2 s	700 ms	How long a mic takes to come up or go down.
Between talkers	–24 to 0 dB	–6 dB	How far the other mics go down when nobody speaks. The mic of the last speaker stays open, so the room noise does not pump. 0 leaves every mic alone.
Speech starts at	+4 to +24 dB	+10 dB	How far above its own mic's noise a voice must be to count as speaking. Lower it for a quiet talker; raise it if breaths open mics.
Remove bleed	Keep / Remove	Keep	Experimental: takes a speaker's voice out of the other mics.
LEAD	per mic	off	Marks a lead mic.
Auto	button		Sets it all for you from your recordings.

CAUTION Remove bleed reads the other tracks while it plays, so on long files it can cause dropouts. Try Keep bleed if you hear any.

TIP If you hear the room noise change when people stop talking, make Between talkers smaller (toward 0 dB): the mics then move less in the pauses.

Specifications

Formats	Audio Unit only (needs ARA)
Hosts	Logic Pro, run with Rosetta
Channels	mono or stereo, one instance per mic
Window	800 × 650, can be resized

How it works

a closer look

KA Auto Mixer reads all the mics at once. It measures each one against the noise of its own mic, and works out for every moment who is speaking: the mic that is clearly loudest, and above Speech starts at over its own noise. When it is a close call, a lead mic wins.

A voice is also heard a little later in the mics of the people next to the speaker. Where two or more mics are loud, KA Auto Mixer measures which one hears the sound first; a mic that hears it later, and quieter, is only hearing bleed, and is not opened.

The mic that is speaking is left where it is. The others are turned down by Turn down others, which takes out most of the bleed. Changes take Fade time to happen, so the mixing does not jump. In the pauses the mic of the last speaker stays open and the others go down by Between talkers.

Auto measures how much each mic leaks into the others and sets Turn down others a little above it, and sets Speech starts at for you.

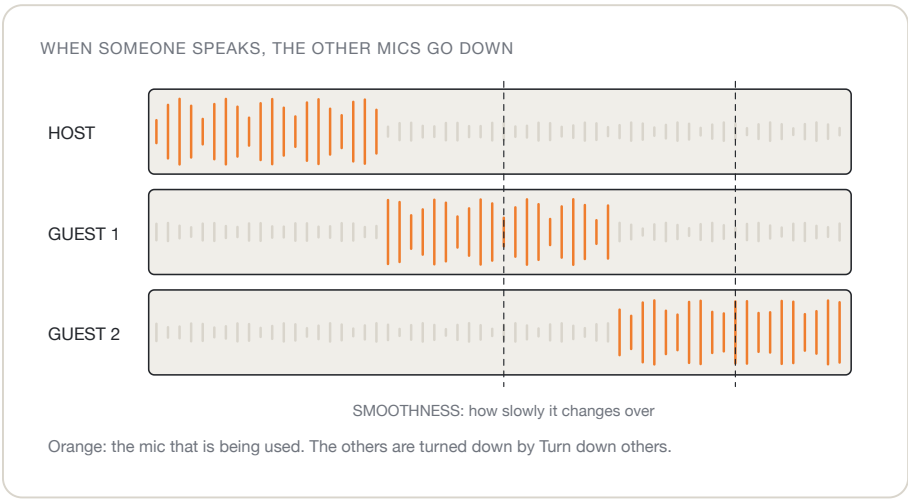


FIG. 26 Whoever speaks is used; the others are turned down.

NO. 11

KA Saturator

Adds warmth and weight by bending the sound: Soft, Warm or Hard.

KA Saturator adds warmth and weight by bending the sound. Drive pushes the sound into a curve, and Soft, Warm and Hard are three different curves. It runs at four times the sample rate so that the new harmonics stay clean, and the loudness stays about the same as you turn Drive up, so that you can judge the tone and not only the level.

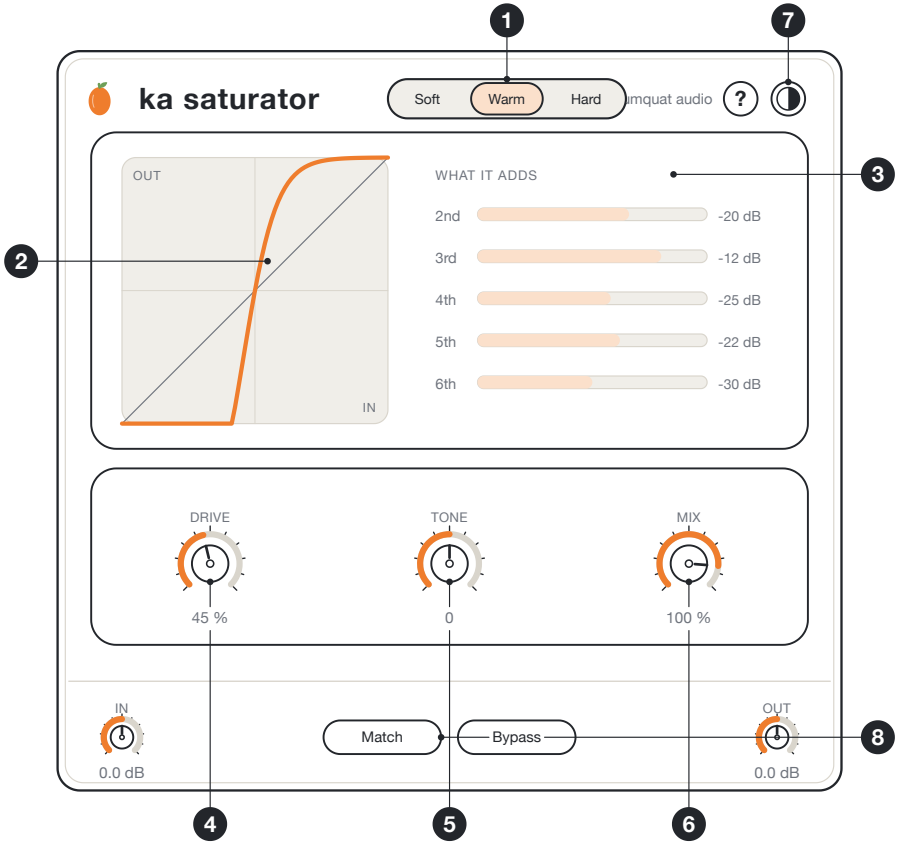


FIG. 27 The parts of KA Saturator. The numbers are explained on the next page.

Parts

fig. 27

1. **Type**
   Soft, Warm or Hard: the shape of the curve that bends the sound.
2. **The curve**
   What the sound is bent by. The straight line is the sound left alone; the bright part is the part your music is using.
3. **What it adds**
   The harmonics that the curve makes, in dB below the main tone. Drag a bar to set that harmonic's level yourself (the small mark shows what the curve makes); double-click a bar to give it back to the curve.
4. **Drive**
   How hard the sound is pushed into the curve. The loudness stays about the same.
5. **Tone**
   Tilts the sound after the saturation: darker to the left, brighter to the right.
6. **Mix**
   How much of the saturated sound is blended with the original.
7. **Help and light/dark**
   The ? turns the pop-up help on and off. The half circle switches light and dark.
8. **In, Out and Match**
   In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

Operation

step by step

1. Put the plug-in on a track or a bus and play something with some body in it: drums, bass or a voice.
2. Choose a Type. Soft rounds the peaks gently. Warm is thicker and rounder. Hard keeps the sound clean until a ceiling and then cuts off flat: edgy and dense.
3. Raise Drive until you hear the tone change. The picture shows the curve bending, and the harmonics it adds.
4. Use Tone to tilt the result darker or brighter, and Mix to blend it with the original (parallel saturation).
5. Switch the plug-in off and on to compare. The level is already matched, so what you hear change is the tone. Use Out to trim.
6. To shape the harmonics yourself, drag a bar in "What it adds" to the level you want for that harmonic (the small mark shows what the curve makes). Double-click the bar to give it back to the curve. A bar you have set stays where you put it as Drive changes.

Controls and settings

range · start

Control	Range	Starts at	What it does
Type	Soft / Warm / Hard	Soft	The shape of the curve.
Drive	0 to 100 %	30 %	How hard the sound is pushed into the curve. (100 % is 30 dB.)
Tone	–100 to +100 %	0	Tilts the sound after the saturation: down is darker, up is brighter.
2nd – 6th harmonic	off to –0 dB below the tone	as the curve makes them	Drag a bar to set that harmonic's level yourself. Double-click to give it back to the curve.
Mix	0 to 100 %	100 %	How much of the saturated sound is blended with the original.
In / Out	In ±24, Out ±12 dB	0 dB	The two knobs under the meters: before and after the plug-in.
Match	on / off	off	Keeps the output as loud as the input. See In, Out and Gain Match.
Live	on / off	off	No oversampling, so no delay at all: for playing through the plug-in. The curve is the same; very bright sounds at high Drive can fold a little more noise back in.

TIP For playing through the plug-in, switch Live on: the delay is 0 samples. Switch it off again for mixing.

How it works

a closer look

Saturation bends the sound. Every sample is sent through a curve: a quiet sound passes almost unchanged, a loud one is bent, and the bending adds new frequencies, the harmonics of whatever is playing. Drive sets how far into the curve the sound is pushed.

Soft rounds the peaks more and more as they get louder, and adds odd harmonics (the 3rd, the 5th). Warm bends the top and the bottom of the wave differently, which adds the even harmonics as well (the 2nd, the 4th): a thicker, rounder sound. Hard leaves the sound alone until it is close to a ceiling and then cuts it off flat: edgy and dense.

New harmonics can be higher than the highest frequency the system can hold, and then they fold back down as harsh, out-of-tune noise (aliasing). To avoid it the plug-in works at four times the sample rate and works out the curve in a more careful way. The cost is a delay of 61 samples, which your music software is told about.

The loudness is kept about the same as Drive goes up, so that a more saturated sound does not seem better only because it is louder. Tone is a tilt round 650 Hz applied after the saturation. A gentle high-pass takes out the offset that the lopsided Warm curve makes.

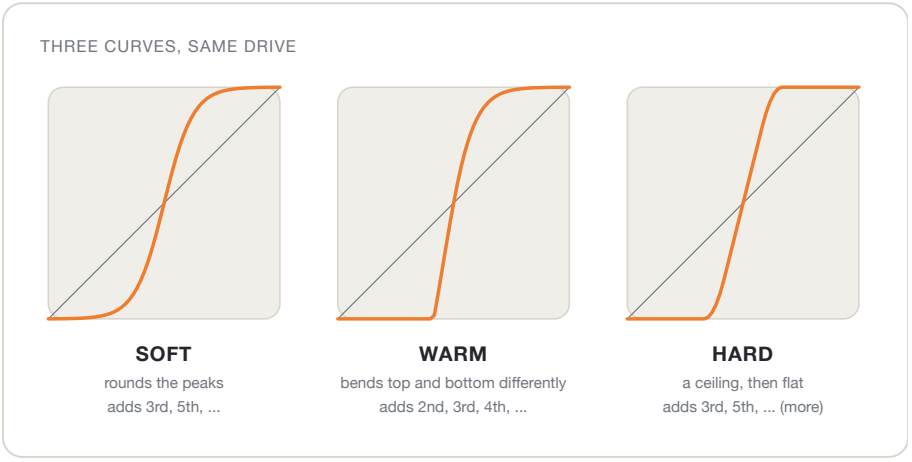


FIG. 28 The three curves, and the harmonics each one adds.

NO. 12

KA Delay

Repeats the sound, in milliseconds or on the beat, with four ways of using the two sides.

KA Delay repeats the sound. Set the time in milliseconds, or as a note of the song's tempo so that the repeats stay on the beat. Feedback sends the repeats round again, Tone makes each one darker or thinner, High-pass and Low-pass take the low end or the top out of them, and Stereo chooses how they move between the two sides.

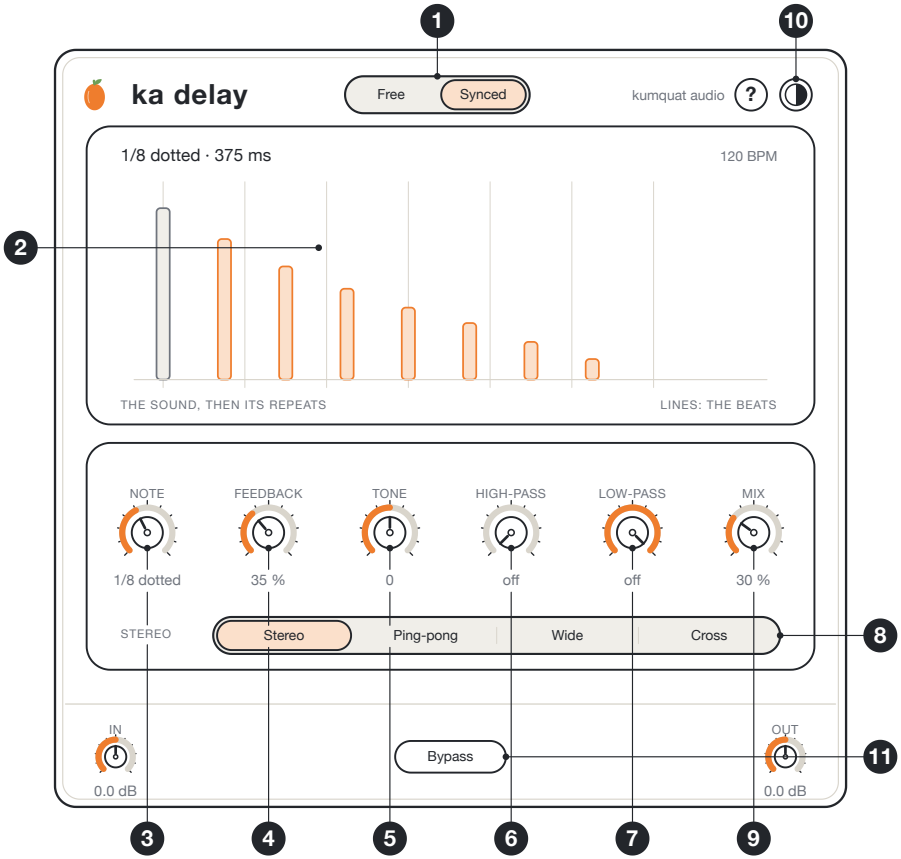


FIG. 29 The parts of KA Delay. The numbers are explained on the next page.

Parts

fig. 29

1. **Free / Synced**
   Free: the time is in milliseconds.
   Synced: the time is a note of the song's tempo.
2. **The repeats**
   The tall bar is the sound; the bars after it are its repeats, where they fall in time and how loud each one is. The thin lines are the beats.
3. **Time or Note**
   The time between repeats: milliseconds when Free, a note value when Synced.
4. **Feedback**
   How much of each repeat goes round again.
5. **Tone**
   Each repeat is darker (left) or thinner (right) than the one before.
6. **High-pass**
   Takes the low end out of the repeats, a little more with every one, so they stay out of the way of the bass. Off at the left.
7. **Low-pass**
   Takes the top out of the repeats, a little more with every one. Off at the right.
8. **Stereo**
   How the repeats use the two sides: Stereo, Ping-pong, Wide or Cross.
9. **Mix**
   How much echo is mixed in with the original.
10. **Help and light/dark**
    The ? turns the pop-up help on and off. The half circle switches light and dark.
11. **In and Out**
    In turns the sound up or down before the plug-in hears it, and Out turns the result up or down.

Operation

step by step

1. Put the plug-in on a track, or better on an effects bus that you send to.
2. Choose Free to set the time in milliseconds, or Synced to choose a note. 1/8 dotted is the classic.
3. Set Feedback: low gives one or two repeats, higher gives a trail of them.
4. Turn Tone down to make each repeat darker, so that they sit behind the sound, or up to make them thinner, so that they stay out of the way.
5. Use High-pass to take the low end out of the repeats (so they stay out of the way of the bass) and Low-pass to take the top out (a tape or analog delay): each repeat goes through them again, so they get stronger with every repeat.
6. Choose how the repeats use the two sides with Stereo, and set Mix.

Controls and settings

range · start

Control	Range	Starts at	What it does
Free / Synced	two choices	Synced	Time in milliseconds, or a note of the song's tempo.
Time	1 ms to 2 s	350 ms	The time between repeats (Free).
Note	1/32 to 1 bar	1/8	The time between repeats (Synced), with triplets and dotted notes.
Feedback	0 to 100 %	35 %	How much of each repeat goes round again.
Tone	–100 to +100 %	0	Down: each repeat is darker. Up: each is thinner.
High-pass	20 Hz to 2 kHz	off	Takes the low end out of the repeats (12 dB an octave). Off at the left.
Low-pass	500 Hz to 20 kHz	off	Takes the top out of the repeats (12 dB an octave). Off at the right.
Stereo	Stereo / Ping-pong / Wide / Cross	Stereo	How the repeats use the two sides.
Mix	0 to 100 %	30 %	How much echo is mixed in with the original.
In / Out	–24 to +24 dB	0 dB	The two knobs under the meters: before and after the plug-in.

TIP When the delay is synced, the thin lines in the picture are the beats of the song: the repeats sit on them.

TIP Changing Time while the music plays makes the repeats slide in pitch a little, like a tape delay. Leave it alone for the cleanest sound.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo
Delay	none
Longest time	6 seconds
Window	700 × 520, can be resized

How it works

a closer look

The sound goes into a delay line and comes out Time later. Part of what comes out goes back in (Feedback), so it comes round again, quieter, and again. The repeats are as far apart as the time you set.

Tone is a filter in that loop. Every trip round the loop goes through it, so each repeat is a little darker (or thinner) than the one before. A soft limiter in the loop keeps it from running away, so even Feedback at 100 % dies away.

Stereo: each side has its own repeats. Ping-pong adds the two sides together and sends the first repeat to the left, the next to the right, and so on. Wide makes the left side's time three quarters of the right side's, so the two weave across the picture. Cross keeps the two sides apart but swaps them on every trip round.

When the time is a note, the plug-in asks your music software for the tempo. If the software does not give one, it uses 120 BPM, and the picture says so. The time glides when you change it, so there are no clicks.

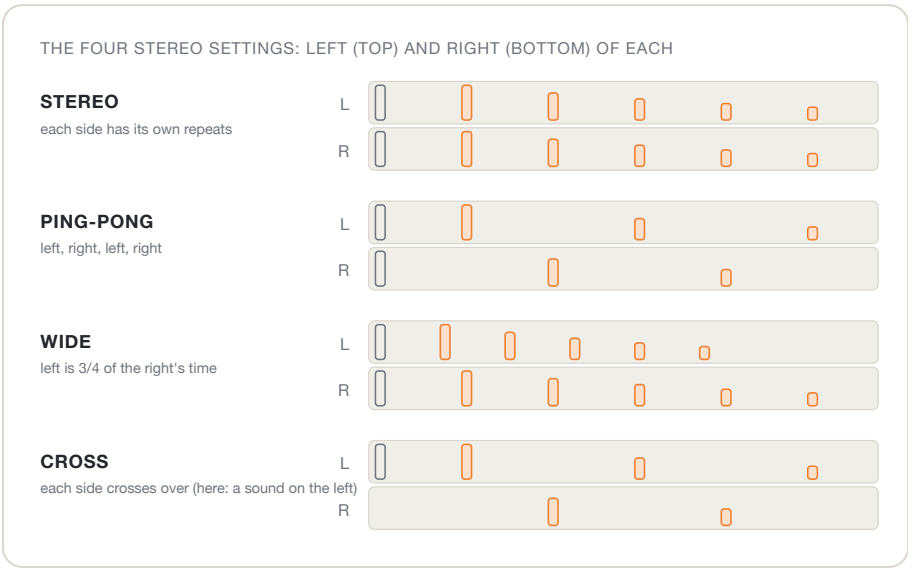


FIG. 30 How the four Stereo settings place the repeats.

NO. 13

KA Bus Comp

A compressor for gluing a mix or a group of tracks together.

KA Bus Comp is a compressor for gluing a mix, or a group of tracks, together. It works on the stereo sound as a whole, it does not listen to the bass (so the kick does not make everything pump), and its Auto release follows the music.

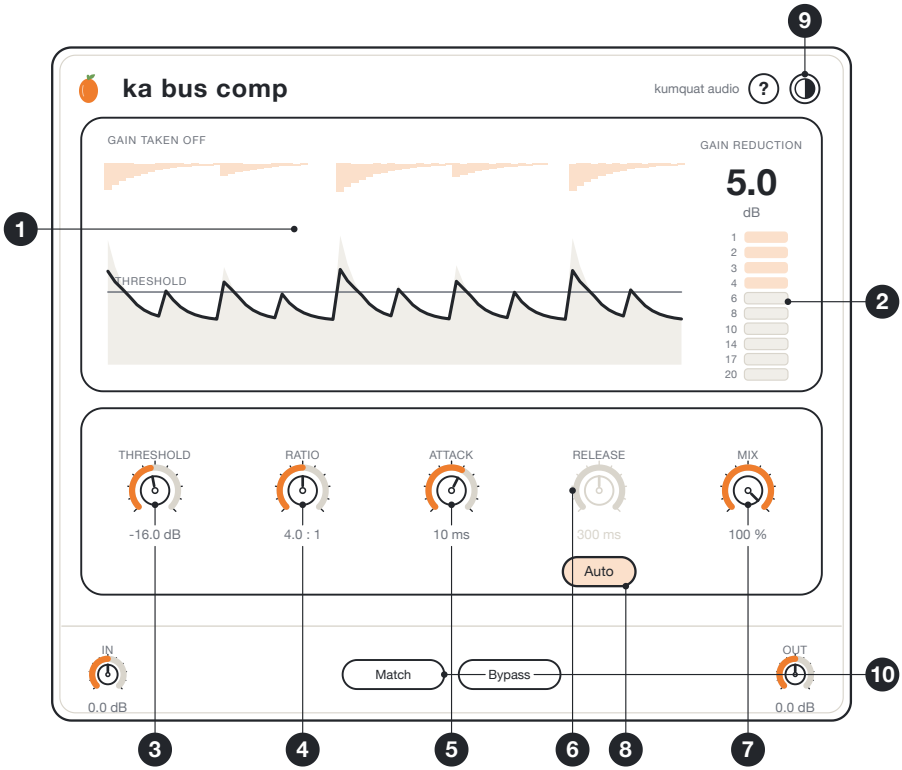


FIG. 31 The parts of KA Bus Comp. The numbers are explained on the next page.

Parts

fig. 31

1. **The picture**
   The level coming in (grey), the level going out (the line), the threshold (a thin line), and the gain being taken off, hanging from the top in orange.
2. **Gain reduction**
   How many dB it is turning the sound down now, on a ladder of ten lights: each comes on when that much is being taken off, and the last three are red. The number holds the deepest value for a moment so that it can be read; the outlined light is the deepest lately.
3. **Threshold**
   The level above which the sound is turned down.
4. **Ratio**
   How firmly it is turned down once it is over the threshold. It clicks between 2, 4 and 10.
5. **Attack**
   How fast it reacts to a loud sound. It clicks between 0.1, 0.3, 1, 3, 10 and 30 ms. Slower keeps the punch of each hit.
6. **Release**
   How fast it lets go: 100 ms, 300 ms, 600 ms or 1.2 s. Not used while Auto is on.
7. **Mix**
   How much compressed sound is blended with the original.
8. **Auto**
   Chooses the release from the music: short after a lone hit, longer when the mix has been held down for a while.
9. **Help and light/dark**
   The ? turns the pop-up help on and off. The half circle switches light and dark.
10. **In, Out and Match**
    In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness. The gain to make up for the compression is the Out slider.

Operation

step by step

1. Put the plug-in on a bus or on the master, and play the loudest part.
2. Lower Threshold until the meter shows 1 to 3 dB of gain reduction on the loud parts.
3. Choose Ratio: around 2 is gentle, 4 is the usual setting for a bus, 10 is firm.
4. Set Attack. A slower attack (10 to 30 ms) lets the start of each hit through, which keeps the punch.
5. Leave Release on Auto, or set it by ear: short follows the rhythm, long is smooth.
6. Use the Out knob (at the bottom right) to bring the level back up, and compare with the plug-in switched off, or press Match. Mix below 100 % is parallel compression.

Controls and settings

range · start

Control	Range	Starts at	What it does
Threshold	–60 to 0 dB	–14 dB	The level above which the sound is turned down.
Ratio	2, 4, 10	4	How firmly it is turned down once it is over the threshold. The knob clicks between the steps.
Attack	0.1, 0.3, 1, 3, 10, 30 ms	10 ms	How fast it reacts to a loud sound. The knob clicks between the steps.
Release	100, 300, 600 ms, 1.2 s	300 ms	How fast it lets go (used when Auto is off). The knob clicks between the steps.
Auto	on / off	on	Chooses the release from the music.
Mix	0 to 100 %	100 %	How much compressed sound is blended with the original.
In / Out	–24 to +24 dB	0 dB	In turns the sound up or down before the plug-in hears it; Out turns what comes out (it is the make-up gain).
Match	on / off	off	Keeps the output as loud as the sound that came in, whatever In and the plug-in do. See In, Out and Gain Match.

TIP If it takes the life out of the drums, make Attack slower or Mix lower.

CAUTION A very fast attack with a very short release makes the compressor follow the waves themselves, which can sound distorted. Slow one of them down.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo (the two sides are linked)
Delay	none
Window	700 × 520, can be resized

How it works

a closer look

The compressor listens to the louder of the two sides, after a gentle high-pass at 90 Hz, and measures how far over the Threshold it is. Over the threshold it turns the sound down, and Ratio says by how much: at 4:1, a sound 8 dB over comes out 2 dB over.

The knee is soft: the compressor starts to act a little before the threshold and reaches its full ratio a little after it, so it comes in smoothly.

Attack and Release say how quickly the gain falls and recovers. Auto release follows the music: short after a lone hit, up to 1.2 seconds when the mix has been held down for a while.

Mix blends the compressed sound with the original, for parallel compression. There is no make-up knob: the Out knob at the bottom right is the make-up gain.

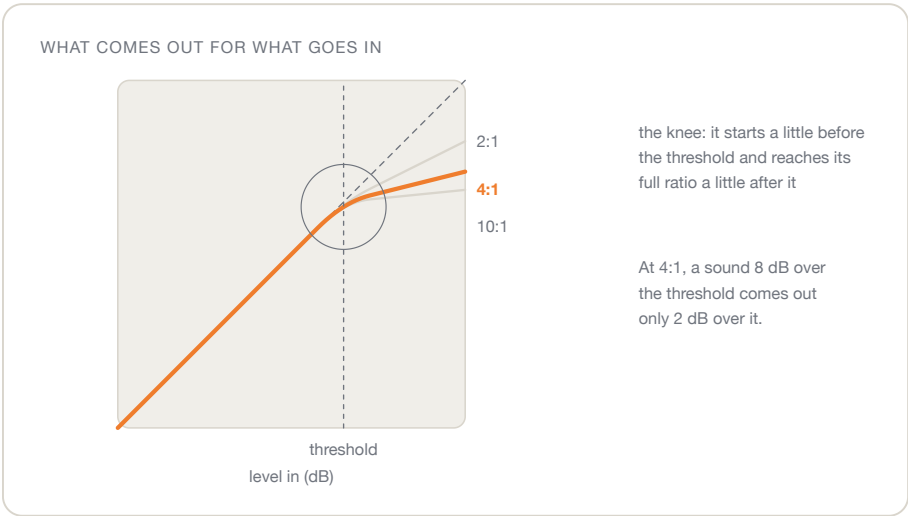


FIG. 32 The curve of the compressor, with its soft knee.

NO. 14

KA Reverb

A reverb with five spaces (one plays a room from a sound file) and a send EQ. A drum kit sounds like it is in the room.

KA Reverb is a reverb for putting sound in a space. Choose the kind of space (Room, Chamber, Hall, Plate or Room file), set its Size and Decay, and shape the tail with Bass, Highs and Mod. The Send EQ filters what goes into the reverb. The first reflections follow the walls of a room of the size you set, so a drum kit sent through it sounds like it was recorded there.

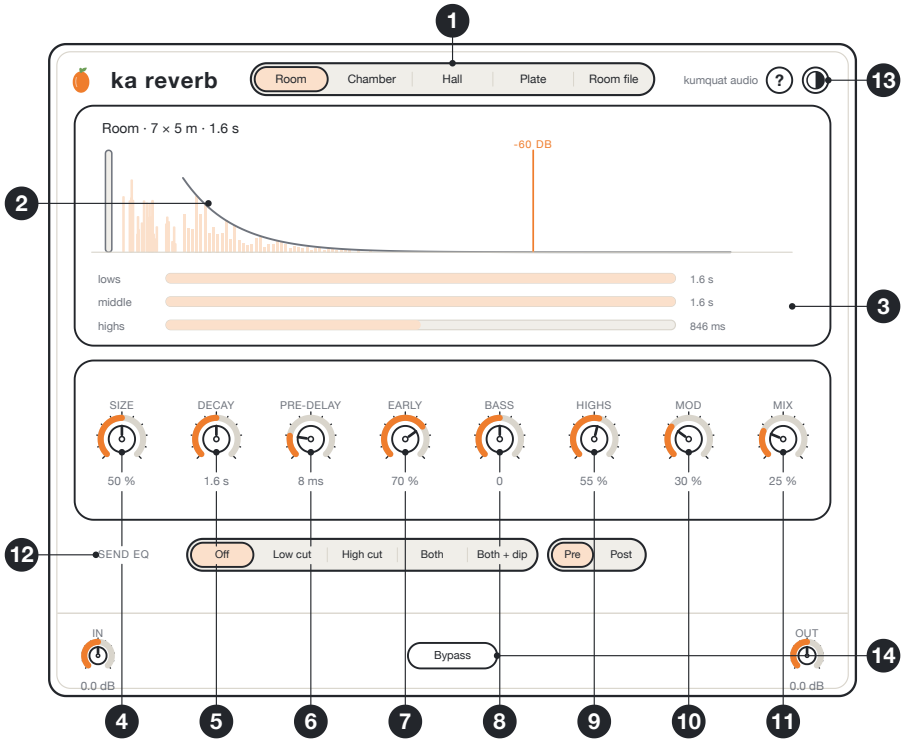


FIG. 33 The parts of KA Reverb. The numbers are explained on the next page.

Parts

fig. 33

1. **Space**
   Room: the first reflections follow the walls of a room. Chamber: bigger, smoother. Hall: large and slow. Plate: no room, a dense shimmer. Room file: a room loaded from a sound file (Load room... appears on the picture).
2. **The space**
   What a clap would sound like here: the grey bar is the sound, the tall bars the first reflections off the walls, then the tail. The thin line: the time to fall by 60 dB.
3. **Lows, middle, highs**
   How long each part of the sound lasts. Bass and Highs move them.
4. **Size**
   The size of the space: how far away the walls are, so how soon the first reflections come.
5. **Decay**
   How long the reverb lasts: the time it takes to fall by 60 dB.
6. **Pre-delay**
   The wait before the reverb starts.
7. **Early**
   How loud the first reflections are. Up: the sound sits in the room. Down: mostly tail, and further away.
8. **Bass**
   How long the lows last compared with the middle. Up: they bloom. Down: the reverb stays clean and tight.
9. **Highs**
   How soon the highs are lost. Down: dark, soft walls. Up: bright and airy.
10. **Mod**
    How much the tail drifts: lush and alive, and never ringing at one pitch.
11. **Mix**
    How much reverb is mixed in with the original.
12. **Send EQ**
    Filters the reverb, like an EQ on an aux send. Off: nothing. Low cut: high-pass (no rumble). High cut: low-pass (soft top). Both: high-pass and low-pass. Both + dip: both and a gentle scoop around 800 Hz, to keep the reverb out of the voice. Each choice is drawn with the shape of what it lets through. Pre: it filters what goes into the reverb. Post: it filters the reverb that comes out. The dry sound is never touched.
13. **Help and light/dark**
    The ? turns the pop-up help on and off. The half circle switches light and dark.
14. **In and Out**
    In turns the sound up or down before the plug-in hears it, and Out turns the result up or down.

Operation

step by step

1. Put the plug-in on an effects bus (or on a track) and set Mix to taste.
2. Choose a Space: Room, Chamber, Hall, Plate or Room file. Set Size (the walls move away as you turn it up) and Decay.
3. Use Early for how much of the room you hear: up puts the sound inside it, down moves it back.
4. Shape the tail with Bass, Highs and Mod. The bars in the picture show how long the lows, the middle and the highs last.
5. Add Pre-delay, 10 to 40 ms for instance, to keep a voice clear in front of its reverb.

Controls and settings

range · start

Control	Range	Starts at	What it does
Space	Room, Chamber, Hall, Plate, Room file	Room	The kind of space. Room file plays a room loaded from a sound file (Load room...).
Size	0 to 100 %	50 %	The size of the space, from small to very large.
Decay	0.2 to 10 s	1.6 s	How long the reverb lasts (the time to fall by 60 dB).
Pre-delay	0 to 250 ms	8 ms	The wait before the reverb starts.
Early	0 to 100 %	70 %	How loud the first reflections are. In Room file: how loud the first 80 ms of the loaded room is against the rest.
Send EQ	Off, Low cut, High cut, Both, Both + dip	Off	Filters the reverb. Low cut: high-pass at 150 Hz. High cut: low-pass at 5.5 kHz. Both: both. Both + dip: both, and a gentle scoop near 800 Hz.
Send EQ position	Pre, Post	Pre	Pre: the EQ filters what goes into the reverb. Post: it filters the reverb that comes out.
Bass	–100 to +100 %	0	How long the lows last compared with the middle.
Highs	0 to 100 %	55 %	How soon the highs are lost. Down: dark. Up: bright and airy.
Mod	0 to 100 %	30 %	How much the tail drifts. A little keeps it lush and alive.
Mix	0 to 100 %	25 %	How much reverb is mixed in with the original.
In / Out	–24 to +24 dB	0 dB	The two knobs under the meters: before and after the plug-in.

TIP To put a drum kit in a room: use Room on an effects bus at Mix 100 %, and raise the send until the kit sits in it. Early does most of the work; Pre-delay

How it works

a closer look

First the sound waits for the Pre-delay. Then it goes three ways: to the early reflections, into a smear, and from there into the tail.

The early reflections follow the walls of a room. The plug-in takes a box-shaped room (a Room is 6.4 by 4.8 by 3 metres, a Hall 30 by 21 by 13), puts two sources a little apart in it, and two ears, and follows the bounces off every wall, the floor and the ceiling. Each bounce arrives at its own moment, from its own side, and a little quieter and darker than the last.

The smear is a few all-pass filters, so that a snare hit becomes a soft cloud and not a string of separate echoes. The tail is sixteen delay lines whose sound is mixed together on every trip round, so one echo becomes many very quickly. Each line drifts a little (Mod), so that nothing rings at one pitch.

On each trip the sound is made quieter, so that the tail falls by 60 dB in the time set by Decay. Bass makes the lows fall more slowly or more quickly than the middle, and Highs sets how soon the highs are lost. Size moves the walls: the reflections come later and the echoes of the tail are spaced further apart.

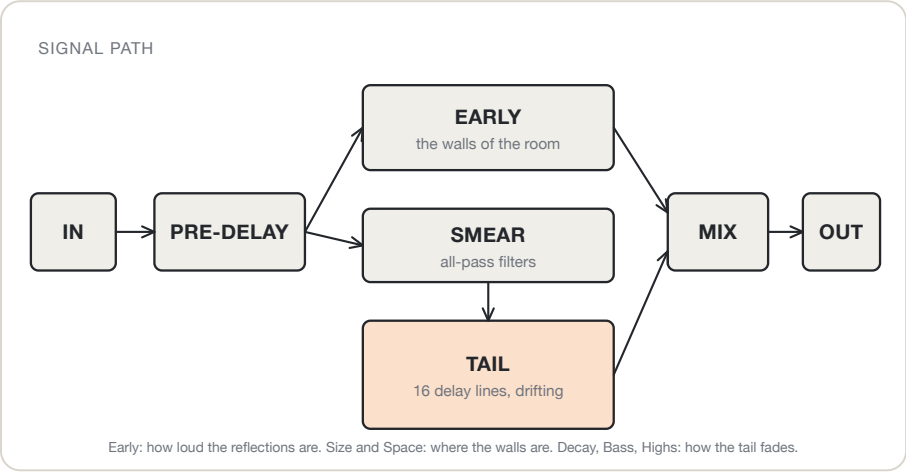


FIG. 34 From the sound to the reverb.

NO. 15

KA Time Align

Lines up several mics on one sound: delays, polarity and phase, with the numbers for your console.

KA Time Align lines up several mics that hear the same sound. It reads the whole recording, works out how much later each mic hears the sound than the reference mic (a mic that was not played, or that hears a different voice, is left alone), and gives you each delay in milliseconds, samples and metres: the numbers to set on a console, or let the plug-in delay the tracks for you. It also finds mics that are upside down, and adds a phase filter where one helps the low end.

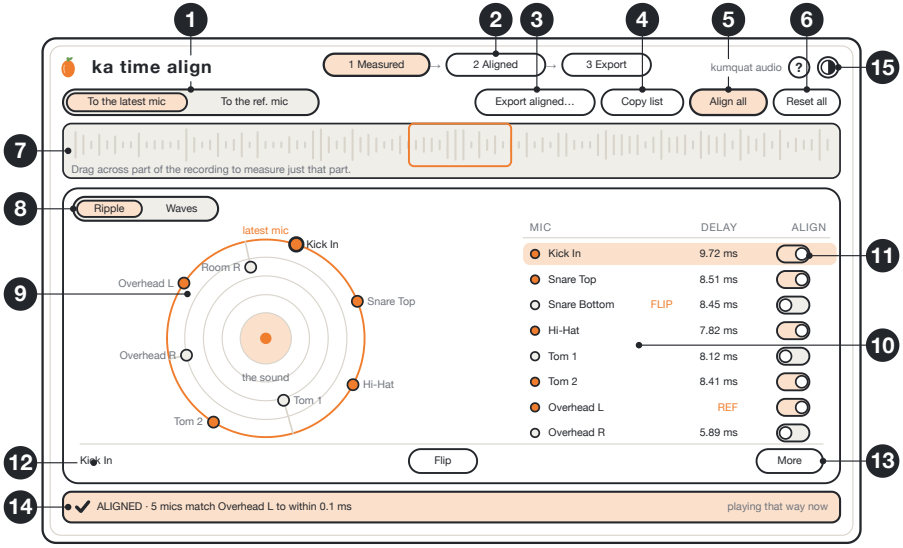


FIG. 35 The parts of KA Time Align. The numbers are explained on the next page.

Parts

fig. 35

1. **Latest / reference**
   To the latest: every delay is zero or more, as a console needs. To the reference: every mic lines up with the reference.
2. **The three steps**
   Measured: the mics have been measured. Aligned: they play lined up. Export: write them out. The next step to take is the one drawn with the filled button.
3. **Export aligned**
   Writes each mic to its own WAV file, aligned, in a new folder.
4. **Copy list**
   Copies mic, ms, samples and metres to the clipboard.
5. **Align all**
   Lines up every mic that was measured: delay, flip and phase filter. A mic that was not played, or that hears a different voice, is left alone.
6. **Reset all**
   Puts every mic back as recorded.
7. **The strip**
   The reference's recording. Drag across a few moments of one sound to measure only those; double-click for the whole.
8. **Ripple / Waves**
   Ripple shows the sound spreading out from the middle. Waves shows one mic's wave on the reference's; there, Before and After switch between how the mics were recorded and how they play now.
9. **The ripple**
   Each mic is a dot. The later a mic hears the sound, the further out it sits. Align slides the dots to one ring, the thin one, and a ripple spreads from the middle.
10. **The list**
    Each mic with the delay to add. Click a name to look at it; double-click to make it the reference (REF). FLIP means the mic is upside down compared with the reference.
11. **Align**
    Plays the mic with its delay, flip and phase filter.
12. **The mic being looked at**
    Flip turns it upside down. A red dot beside a mic's delay means it still does not match the reference.
13. **More**
    Shows the fine-tuning: other moments, other places the mic could be lined up, Flip, – and + and the by-hand number. You do not need them for most sessions.
14. **The bar**
    What the plug-in is doing and whether the mics match. It fills while it works.
15. **Help and light/dark**
    The ? turns the pop-up help on and off. The half circle switches light and dark.

Operation

step by step

1. Open Logic Pro using Rosetta, and put the plug-in first in the Audio FX list on the track of every mic. Leave Logic's own polarity button (Φ) off on these tracks.
2. Select the regions and press play. The bar at the bottom shows the work. The overheads are the reference (REF); double-click another name to use that mic instead.
3. When the bar is done, each mic is a dot in the Ripple view: the later a mic hears the

Controls and settings

range · start

Control	Range	Starts at	What it does
To the latest / reference mic	two choices	latest	Latest: every delay is zero or more, as a console needs. Reference: every mic lines up with the reference and may move earlier.
Align all, Reset all	buttons		Lines up every mic that was measured, or puts every mic back as recorded.
Copy list	button		Copies mic, ms, samples and metres to the clipboard.
Export aligned	button		Writes each mic to its own WAV file, aligned, in a new folder.
The strip	drag	whole recording	Drag across the hits of one drum to measure only those.
Ripple / Waves	two choices	Ripple	Ripple shows each mic as a dot, further out the later it hears the sound. Waves shows one mic's wave on the reference's.
Before / After	two choices	After	In Waves: shows every wave as recorded, or as it plays now.
More	button		Shows the fine-tuning: other moments, other places the mic could be lined up, Flip, – and +, and the by-hand number (drag it up for later, down for earlier).

CAUTION The plug-in reads the recordings before Logic's mixer, so it cannot see Logic's polarity button (Φ), and its own flip can undo it. Switch Φ off on the tracks you align.

TIP Drag across the hits of one drum in the strip, and only those are measured: line up the kick mics from the kick hits, then the snare mics from the snare hits.

TIP A mic that still has a red dot after Align all does not match the reference yet. Press More and try the other places it could be lined up.

Specifications

The list

one row for each mic

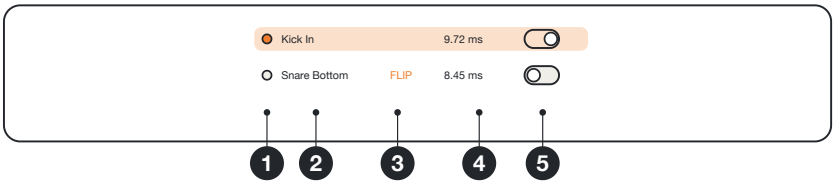


FIG. 36 A close look at two rows of the list.

1. **The dot**
   Orange when the mic is lined up, pale when it is not.
2. **Mic**
   The track's name. Click it to look at the mic; double-click makes it the reference (REF). The orange row is the mic being looked at.
3. **FLIP**
   The mic is upside down compared with the reference. Press Flip to turn it over.
4. **Delay**
   How long to delay this mic so that it lines up, in milliseconds.
5. **Align**
   Plays the mic with its delay, flip and phase filter.

How it works

a closer look

Two mics on one drum hear the same sound at different moments, because sound travels. Added together, the copies cancel in places and the drum sounds thin. The fix is to delay the mic that hears it first by the difference.

KA Time Align measures that difference on the recording itself. It listens to the whole recording from the low end up to the sharp part of each hit, and a stretch counts for more where that mic is the loud one, so the kick mic is judged mostly on the kick.

It also takes each mic's own strong hits, finds the same hit in the reference, and sees how far apart they are. When many hits agree on one delay, that delay is the answer. Mics that agree this way are settled first, and the others are made to add up with them.

A mic can also be upside down. That shows as the two cancelling, and the plug-in offers to flip the mic.

Sometimes a delay lines up the highs but not the lows. A phase (all-pass) filter turns the phase round one pitch without changing how loud anything is. The plug-in keeps one only when it clearly helps; the percentage shows by how much. Two mics on one drum are also checked against each other.

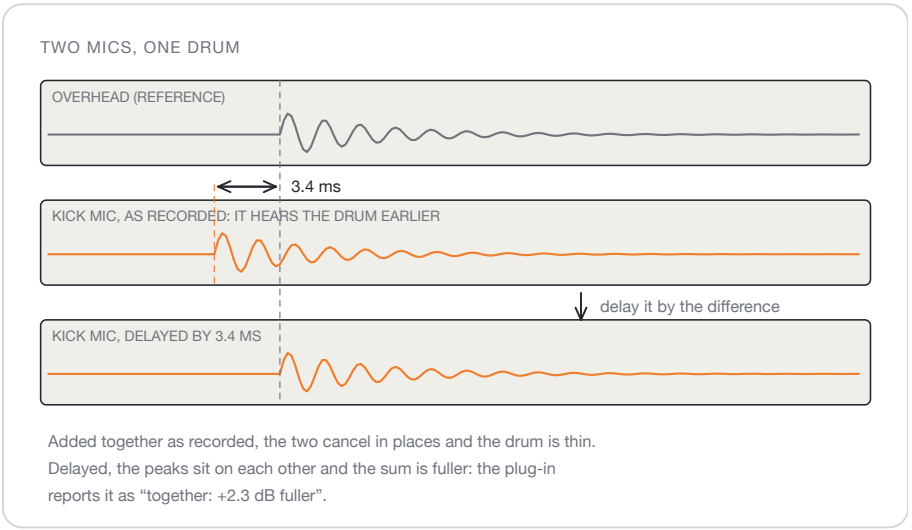


FIG. 37 Delaying the early mic makes the two add up.

NO. 16

KA Multiband Comp

Up to six bands, each with its own compressor or expander, in minimum or linear phase.

KA Multiband Comp splits the sound into up to six bands and gives each band its own compressor or expander. Drag the crossovers on the picture of the spectrum, drag each band's threshold, and see how much each band is being turned down. Use it to tame a boomy low end, a harsh upper mid or a shrill top without touching the rest of the sound.

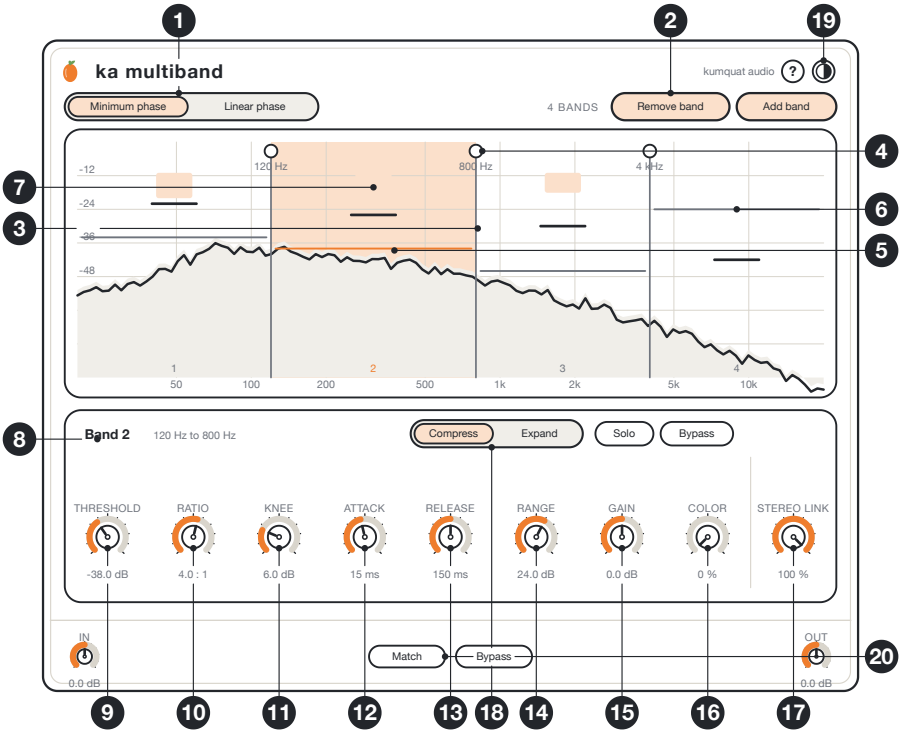


FIG. 38 The parts of KA Multiband Comp. The numbers are explained on the next page.

Parts

fig. 38

1. **Phase**
   Minimum phase: no delay, and the crossovers turn the phase of the sound a little. Linear phase: no phase turning at all, the cleanest on a mix or a master, but the sound is delayed by about 48 ms (the plug-in tells your music software).
2. **Add band / Remove band**
   Splits the selected band in two, or takes it out (its part of the spectrum joins its neighbour). Two to six bands. You can also double-click the picture to add a crossover, and double-click a crossover to remove it.
3. **The spectrum**
   Grey is the sound coming in, the line is the sound going out. Click a band to choose it: the one being edited is shaded.
4. **A crossover**
   Drag the round handle left or right to move the crossover. The frequency is written under it.
5. **The threshold**
   The line across a band: drag it up and down. The compressor turns the band down when its level is above the line (in Expand: below it).
6. **The band's level**
   The short dark bar in each band is the band's own level, to compare with the threshold.
7. **Gain reduction**
   The orange block hanging from the top is how much the band is being turned down now, in dB.
8. **The band**
   Which band the controls below are for, and the frequencies it covers.
9. **Threshold**
   The band's level above which it is turned down (in Expand: below which).
10. **Ratio**
    How firmly. 2 is gentle, 4 is firm, 10 and
11. **Knee**
    How gradually the turning down starts round the threshold. 0 is a hard corner.
12. **Attack**
    How fast the band is turned down when it gets loud.
13. **Release**
    How fast the band comes back up.
14. **Range**
    The most the band is ever turned down. A small range makes a gentle, safe correction.
15. **Gain**
    Turns the band up or down by a fixed amount after the compressor, like a broad EQ.
16. **Color**
    A little saturation on this band, after the compressor and its Gain (the same warmth as KA FET Comp and KA Opto Comp). 0 leaves the band clean.
17. **Stereo link**
    How much the two sides are turned down together. 100 % keeps the stereo picture steady.
18. **Compress / Expand, Solo, Bypass**
    Compress turns the band down when it is loud. Expand turns it down when it is quiet, which cleans up rumble and bleed. Solo plays the band alone. Bypass switches this band's compressor off.
19. **Help and light/dark**
    The ? turns the pop-up help on and off. The half circle switches light and dark.
20. **In, Out and Match**
    In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

Controls and settings

range · start

Control	Range	Starts at	What it does
Phase	Minimum, Linear	Minimum	How the bands are split. Linear has no phase turning but delays the sound about 48 ms.
Live	on / off	off	Switches linear phase off (the bands use minimum phase, whatever Phase says), so there is no delay: for playing through the plug-in.
Bands / Crossovers	2 to 6, 20 Hz to 20 kHz	4 bands	Drag the handles on the picture; Add band and Remove band change the number.
Threshold	–60 to 0 dB	–24 dB	The band's level above which it is turned down (Expand: below which).
Ratio	1 to 20	3	How firmly the band is turned down.
Knee	0 to 24 dB	6 dB	How gradually the turning down starts round the threshold.
Attack	0.1 to 250 ms	15 ms	How fast the band is turned down.
Release	10 ms to 2 s	150 ms	How fast the band comes back up.
Range	0 to 40 dB	24 dB	The most the band is ever turned down.
Gain	–24 to +24 dB	0 dB	Turns the band up or down after the compressor.
Color	0 to 100 %	0 %	A little saturation on this band, after the compressor and its Gain. 0 leaves the band clean.
Compress / Expand	two choices	Compress	Compress turns a loud band down; Expand turns a quiet band down.
Solo, Bypass	per band	off	Solo plays one band alone; Bypass switches its compressor off.
Stereo link	0 to 100 %	100 %	How much the two sides are turned down together.
In / Out	–24 to +24 dB	0 dB	The two knobs under the meters: before and after the plug-in.
Match	on / off	off	Keeps the output as loud as the sound that came in, whatever In and the plug-in do. See In, Out and Gain Match.

How it works

a closer look

KA Multiband Comp splits the sound with crossovers. In Minimum phase they are Linkwitz-Riley filters of the fourth order, no delay, and the bands add up to a flat level, with only the phase turning a little round each crossover. In Linear phase each crossover is a long filter cut so that the bands add up to the sound exactly, only later: the phase does not turn at all, and the plug-in tells your music software that it delays the sound by about 48 milliseconds.

Each band listens to its own level (the average over a few periods of its lowest frequency), and compares it with its threshold. A compressor turns the band down as much as Ratio says when it is over the threshold. An expander turns it down when it is under the threshold. The turning down is never more than Range, and it follows Attack and Release.

The bands are then turned up or down by Gain, given a little saturation if you turn Color up (the same warmth as KA FET Comp and KA Opto Comp, for one part of the spectrum), and added up again. Solo plays one band alone. Stereo link makes the two sides follow the louder one.

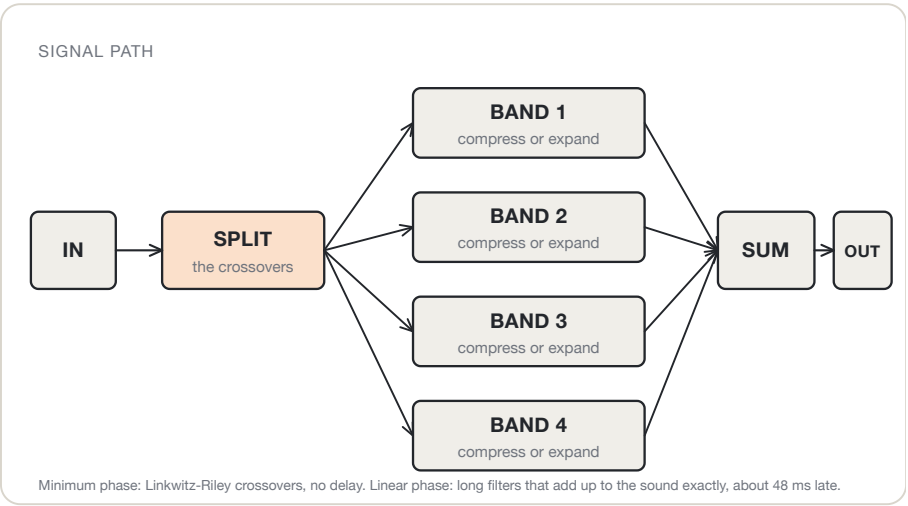


FIG. 39 The sound is split, each band is treated, and the bands are added up.

NO. 17

KA FET Comp

A fast, aggressive compressor that gives drums and vocals their bite.

KA FET Comp is a fast, forward compressor. Its threshold is fixed: you push the sound into it with In and choose how firmly it is held with Ratio. Its attack is very short, so it takes the front off a sound, and its color adds a little saturation the harder it works. Use it on drums, vocals and anything that needs bite.

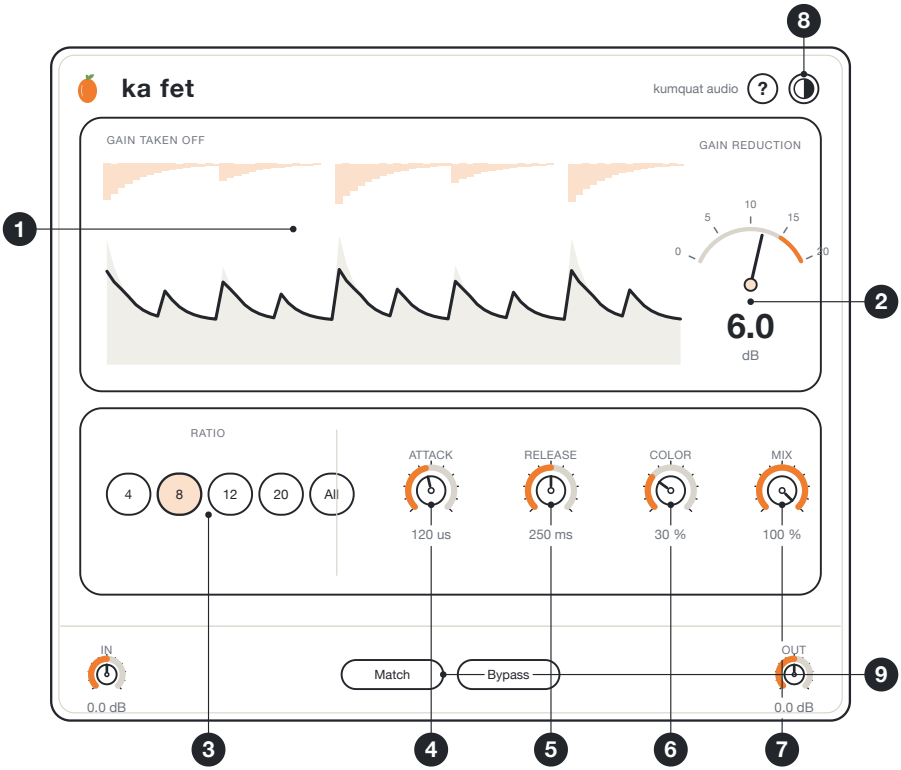


FIG. 40 The parts of KA FET Comp. The numbers are explained on the next page.

Parts

fig. 40

1. **The picture**
   The level coming in (grey), the level going out (the line), and the gain being taken off, hanging from the top in orange. The threshold is fixed.
2. **Gain reduction**
   How many dB it is turning the sound down now, on a needle that flicks with every hit. The last quarter of the dial is red. The number holds the deepest value for a moment so that it can be read.
3. **Ratio**
   How firmly the sound is held once it is over the threshold: 4 to 1 is gentle, 20 to 1 is close to a limiter. All is every button at once: much firmer, a little faster, and more color.
4. **Attack**
   How fast it grabs a loud sound, from 20 microseconds to 0.8 milliseconds. Very fast takes the front off a drum hit; that is where the bite comes from.
5. **Release**
   How fast it lets go. It is quicker after a short peak and up to three times slower after hard work.
6. **Color**
   A little saturation after the compressor, more of it the harder it works. At 0 it is only the compressor.
7. **Mix**
   How much compressed sound is blended with the original (parallel compression).
8. **Help and light/dark**
   The ? turns the pop-up help on and off. The half circle switches light and dark.
9. **In, Out and Match**
   In pushes the sound into the fixed threshold: it is the compressor's Input control. Out is the gain after it. In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

Operation

step by step

1. Put the plug-in on a drum, a voice or a bus and play a loud passage.
2. Turn the In knob (at the bottom left) up until the picture shows the gain reduction you want, 3 to 8 dB for most sources. The threshold is fixed: In is the Input control.
3. Choose a Ratio. 4 is gentle, 8 and 12 are firm, 20 is close to a limiter, and All is every button at once: much firmer and more color.
4. Set Attack. 20 microseconds takes the front off a drum hit and gives it bite; a few hundred lets it through and keeps the punch.
5. Set Release. A fast release is lively and can pump; a slow one is smooth. After hard work it takes up to three times longer. The needle on the right shows each hit.
6. Use Color for a little saturation after the compressor, and Mix below 100 % for parallel compression. Out sets the level you leave with.

Controls and settings

range · start

Control	Range	Starts at	What it does
Ratio	4, 8, 12, 20, All	4	How firmly the sound is held once it is over the fixed threshold.
Attack	20 to 800 µs	120 µs	How fast it grabs a loud sound.
Release	50 ms to 1.1 s	250 ms	How fast it lets go (longer after hard work).
Color	0 to 100 %	30 %	Saturation after the compressor; more of it the harder it works.
Mix	0 to 100 %	100 %	How much compressed sound is blended with the original.
In / Out	–24 to +24 dB	–6 dB / 0 dB	In pushes the sound into the threshold and starts at –6 dB, so it begins gently: turn it up for more. Out is the gain after the compressor.
Match	on / off	off	Keeps the output as loud as the sound that came in, whatever In and the plug-in do. See In, Out and Gain Match.

TIP For a drum bus: Ratio 4, Attack 800 µs, Release 150 ms, and push In for 3 to 4 dB of reduction. For a room-like pump, try All.

TIP The harder it works, the more the top comes forward (up to about 2 dB above 3 kHz) and the sooner the grit starts.

CAUTION A very fast attack with a very short release can sound distorted on low sounds, because the compressor then follows the waves themselves. Slow one of them down.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo (the two sides are linked)
Delay	none
Window	700 × 520, can be resized

How it works

a closer look

The compressor listens to the louder of the two sides with a very fast peak detector, and measures how far the sound is over a fixed threshold. The harder you push the sound in with In, the further over it goes. Ratio says how firmly it is turned down: at 4 to 1, a sound 12 dB over comes out 3 dB over.

The attack is between 20 microseconds and 0.8 milliseconds, so the compressor takes the very front off a sound. The release is program dependent: it is as set after a short peak and up to three times longer after hard work.

All is every ratio button at once. It holds the sound much more firmly, a little faster, with a threshold 4 dB lower, and the color is driven harder. Color is a lopsided saturation after the gain stage, without aliasing, and more of it the harder the compressor is working.

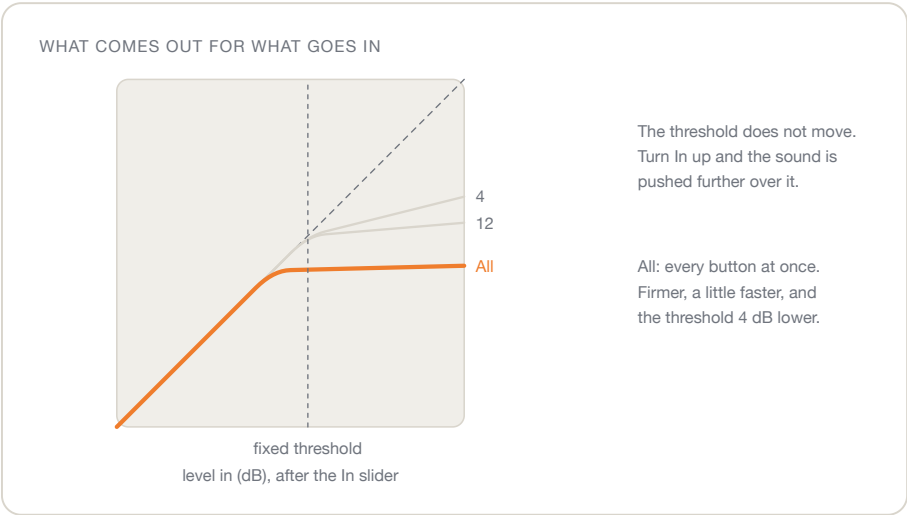


FIG. 41 The threshold is fixed; In pushes the sound over it.

NO. 18

KA Opto Comp

A slow, smooth compressor that levels without being noticed.

KA Opto Comp is a slow, smooth compressor. Peak Reduction sets how much of the sound is held, Compress and Limit choose how firmly, and the release is two-stage and remembers how long it has been working, so it levels a voice or a bass without being noticed.

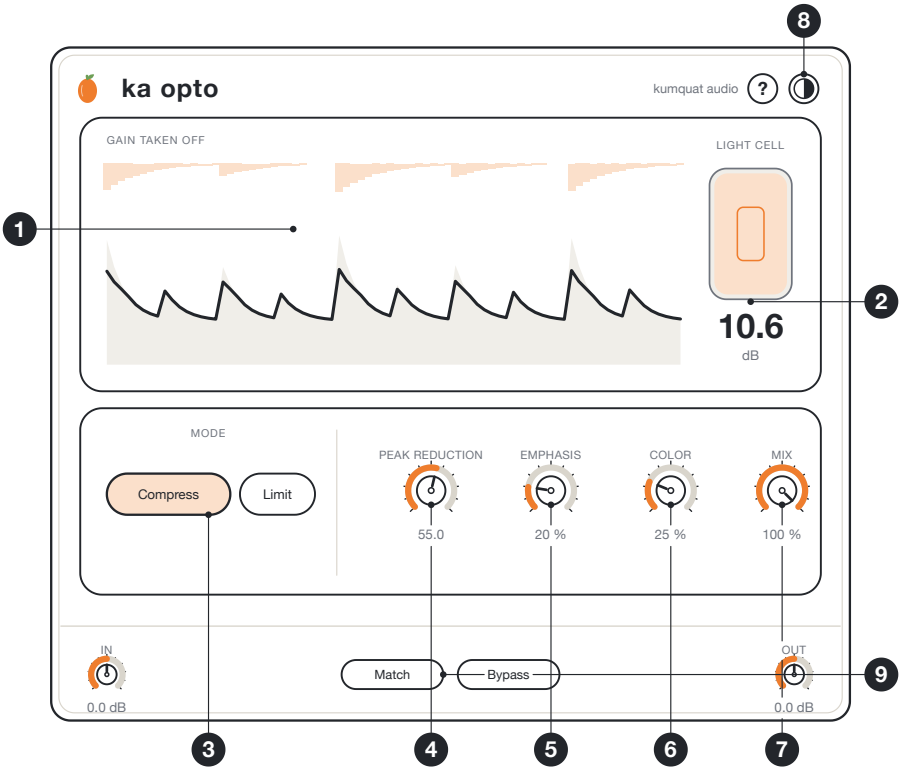


FIG. 42 The parts of KA Opto Comp. The numbers are explained on the next page.

Parts

fig. 42

1. **The picture**
   The level coming in (grey), the level going out (the line), and the gain being taken off, hanging from the top in orange.
2. **Light cell**
   How hard it is working. The bright core follows the sound quickly; the soft halo around it fades slowly. The number is the gain reduction in dB, held at the deepest value for a moment so that it can be read.
3. **Mode**
   Compress: a gentle ratio of about 3 to 1 with a very wide knee, for smoothing and leveling. Limit: about 10 to 1, firm, for catching peaks.
4. **Peak reduction**
   How much of the sound is held: it is the threshold, so more reduction means a lower threshold and more of the sound compressed.
5. **Emphasis**
   Lifts the highs in what the compressor listens to, so sibilance and cymbals are held down more firmly than the rest.
6. **Color**
   A warm saturation after the compressor, more of it the harder it works. At 0 it is only the compressor.
7. **Mix**
   How much compressed sound is blended with the original (parallel compression).
8. **Help and light/dark**
   The ? turns the pop-up help on and off. The half circle switches light and dark.
9. **In, Out and Match**
   In pushes the sound in. Out is the gain after the compressor, which is its make-up gain. In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

Operation

step by step

1. Put the plug-in on a voice, a bass or a bus and play a passage.
2. Turn Peak Reduction up until the picture shows 3 to 6 dB of gain reduction on the loud parts. It is the threshold: more reduction, lower threshold.
3. Choose Compress (gentle, about 3 to 1) or Limit (firm, about 10 to 1).
4. Use Out (the knob at the bottom right) to make up the level. Turn Match on to compare with the plug-in off at the same loudness.
5. Add Emphasis to hold the highs more firmly (sibilance, cymbals). Add Color for a warm saturation, and lower Mix for parallel compression.

Controls and settings

range · start

Control	Range	Starts at	What it does
Peak reduction	0 to 100	28	How much of the sound is held: it is the threshold.
Mode	Compress, Limit	Compress	About 3 to 1 with a very wide knee, or about 10 to 1.
Emphasis	0 to 100 %	0 %	Lifts the highs in what the compressor listens to.
Color	0 to 100 %	25 %	A warm saturation after the compressor; more of it the harder it works.
Mix	0 to 100 %	100 %	How much compressed sound is blended with the original.
In / Out	–24 to +24 dB	0 dB	In pushes the sound in; Out is the gain after the compressor (its make-up gain).
Match	on / off	off	Keeps the output as loud as the sound that came in, whatever In and the plug-in do. See In, Out and Gain Match.

TIP The release is slow. If you need it quicker, use KA FET Comp or KA Bus Comp; if you need it slower and smoother, use KA Opto Comp and let it level.

TIP On a voice, Emphasis around 30 to 60 % makes Opto hold sibilance down a little more.

TIP The harder it works, the softer the very top gets (about 1 dB at 8 kHz, more above) and the warmer its color.

TIP On the light cell, the bright core follows the sound and the halo around it fades slowly.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo (the two sides are linked)
Delay	none
Window	700 × 520, can be resized

How it works

a closer look

The compressor listens to the average level of the louder side, with the highs lifted by Emphasis. Peak Reduction sets the threshold. In Compress the ratio is about 3 to 1 with a very wide knee (18 dB), so it comes in very gradually; in Limit it is about 10 to 1.

The gain is turned down by two lights. The quick one takes about 10 milliseconds to come on and 60 to go off. The slow one takes about 120 milliseconds to come on, and from half a second to many seconds to go off. Together they make the release: half of it is gone in a moment, and the rest creeps away.

How long the slow light stays depends on how hard and for how long the compressor has been working: After a short squeeze it lets go quickly; after a long one it takes seconds.

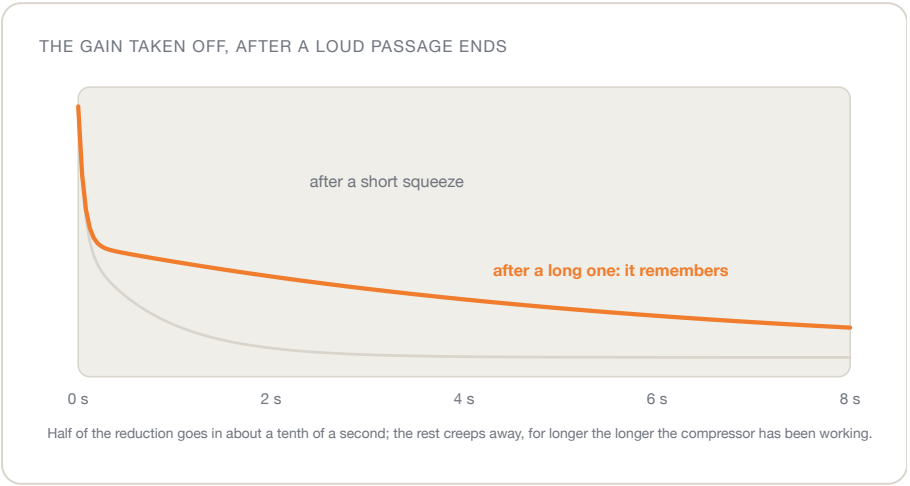


FIG. 43 Half of the reduction goes quickly; the rest creeps away.

NO. 19

KA Rider

Rides the level of a voice for you, around its own level or just above the music.

KA Rider rides the level of a voice the way a person would with a fader: it turns up the quiet words and turns down the loud ones so that the voice sits where you want it, without drawing any automation. It listens to how loud the voice is, learns how loud it usually is, moves the gain towards that level, never by more than Amount, and leaves the gaps between words alone. You can set the level yourself instead, and if you send it the music it keeps the voice just above it.



FIG. 44 The parts of KA Rider. The numbers are explained on the next page.

Parts

fig. 44

1. **The ride**
   The gain it is applying: orange above the line when it turns the voice up, grey below when it turns it down. The thin lines are the limits of Amount.
2. **The voice**
   The voice going in (grey), the voice coming out (the line), and the Level (a thin line). The shaded band at the bottom is Ignore below.
3. **Gain now**
   How many dB it is turning the voice up (+) or down (–) right now, and the same on the meter.
4. **Speed**
   How quickly the gain follows the voice. Slow rides whole phrases; Medium follows words; Fast follows single words. Turning down is always quicker than turning up.
5. **Level**
   Auto: Rider learns how loud the voice usually is and rides around that. Manual: you turn the knob to set it. With a track on the Side Chain input this becomes Over music: the voice is kept that many dB above the music.
6. **Amount**
   How much it rides: the most it will ever turn the voice up or down. A few dB is natural; at 0 it does nothing.
7. **More**
   Opens Ignore below. Anything quieter than it is left alone, so breaths, noise and gaps are not pulled up.
8. **Help and light/dark**
   The ? turns the pop-up help on and off. The half circle switches light and dark.
9. **In, Out and Delta**
   In turns the voice up or down before it is read, and Out turns the result. There is no Match: it would undo the ride. Delta, at the top, plays only what the ride added or took away.

Operation

step by step

1. Put the plug-in on a voice, first in the list if you can, and play a passage with loud and quiet words.
2. Leave Level on Auto. After a few seconds the thin line settles on the voice's own level. Touch a knob and its part of the picture lights up.
3. Set Amount to how much it may move. Start with 4 to 6 dB: a natural ride. Raise it only if the voice is very uneven.
4. Choose Speed. Slow for a smooth ride that cannot be heard; Medium for words; Fast for a very uneven voice.
5. If noise or breaths come up between phrases, press More and raise Ignore below until the gaps stay still. Press Delta to hear only the ride.
6. To keep the voice above the music, send the track to the Side Chain input. Over music then takes the place of Level: turn it to how far above the music you want the voice.

Controls and settings

range · start

Control	Range	Starts at	What it does
Level	Auto, Manual	Auto	Auto learns the voice's level. Manual: you set it.
Amount	0 to 12 dB	6 dB	The most it will turn the voice up or down.
Speed	Slow, Medium, Fast	Medium	How quickly the gain follows the voice.
Ignore below	–80 to –30 dB	–55 dB	Quieter than this, the gain stays where it was.
Over music	–12 to +24 dB	+6 dB	With a track on the Side Chain input, in place of Level: how far above it the voice is kept.
Delta	on / off	off	Plays only what the ride added or took away.
Side Chain: HP and Listen	20 Hz to 1 kHz	off	HP takes the low end out of the music it listens to; Listen plays it.
In / Out	–24 to +24 dB	0 dB	In pushes the voice in; Out is the level you leave with.

TIP Rider works best before the compressor: even out the voice first, then compress what is left.

TIP A fast ride can be heard as the voice breathing. If you hear it, slow the Speed or make Amount smaller.

CAUTION Rider does not know a word from a cough: a loud sound that is not the voice is turned down too.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo (the two sides are linked)
Delay	none
Window	700 × 520, can be resized

How it works

a closer look

KA Rider reads the voice with a short average of its level, after taking out the rumble and the plosives below 100 Hz. It compares that level with the Level and moves the gain by the difference, never by more than Amount.

With Level on Auto, Rider keeps a running average of how loud the voice is while it sings, settles on it in the first few seconds, then follows it slowly. A quiet verse and a loud chorus are each ridden around their own level.

The gain moves smoothly. Turning down is quicker than turning up, so a loud word is caught but a quiet breath is not pulled up. When a loud word starts, the gain turns down at once, so its start is not boosted. Speed sets how quick the rest is.

A second, faster look at the voice also tells it when the voice has stopped or is dying away. Then the gain is not turned up into the silence: it stays where it was, and only turning down is allowed. With a track on the Side Chain input, the Level is replaced by the level of that track plus Over music, as long as the music is playing.

“Loud to quiet 15 → 8 dB” in the picture is the gap between the loud and the quiet words going in and coming out: bring the second number down. A faster Speed and a bigger Amount close the gap more, but they also lift breaths; if they come up, raise Ignore below under More.

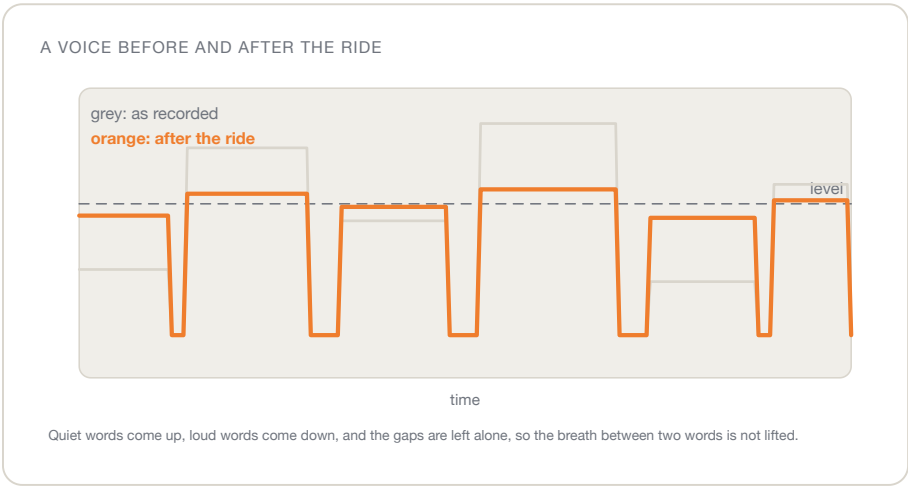


FIG. 45 Quiet words come up and loud ones come down, to the level.

NO. 20 · INSTRUMENT

KA Chladni

A sand-ripple instrument: bring in a sound and play it, with rhythm and harmony.

KA Chladni is an instrument. Every note you play sends a ripple through a plate of sand, and the sand is the picture of the sound. Under it is a simple synth (a wave with something that rings on top), and you can bring in any sound or recording, a World, which is played in tune by the keys: as a pad, or stretched out to a thousand times slower. Hold a note and Pulse plays it again in a rhythm, with harmony. It is made so that you can load a sound and be playing something in seconds, with a few drawings to change it and no menus to find.

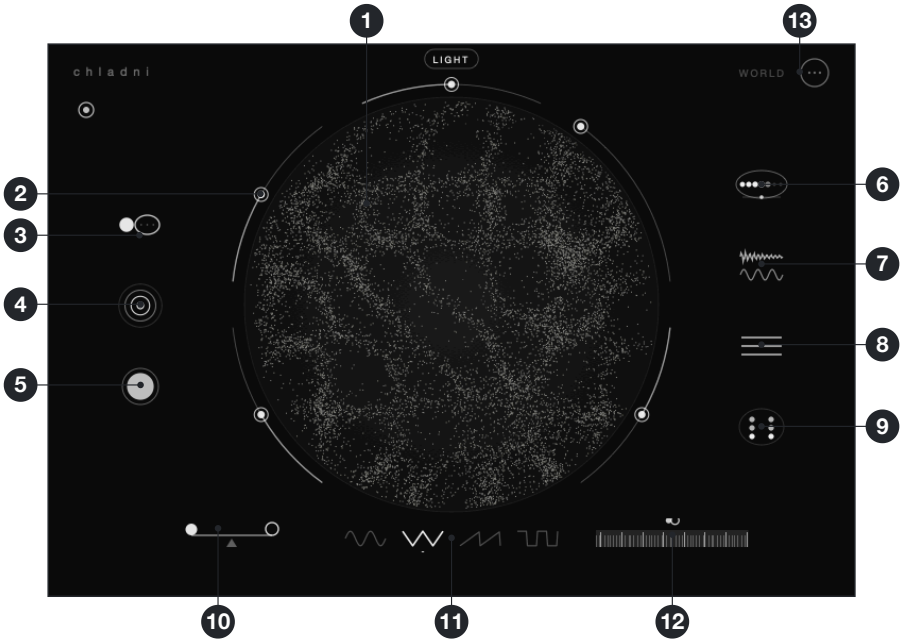


FIG. 46 The parts of KA Chladni. The numbers are explained on the next page.

Parts

fig. 46

1. **The sand**
   Every note you play sends a ripple through the sand. It is the picture of the sound: louder notes move more sand, and the sand settles back to an even spread when the sound stops.
2. **The rim**
   Five arcs round the plate: Bloom (how slowly a note fades in), Pluck (how much a note dies away by itself), Light (how bright), Drive (a little saturation) and Linger (how long it rings after you let go). Each arc has one setting for the synth and one for the world; click the lock that appears beside an arc to move the two together. Right-click an arc to give it a MIDI control.
3. **World**
   Drag up for more of the sound you dropped in. Drag sideways to stretch it, from its own speed to a thousand times slower: a smooth, endless cloud that still follows the pitch of the keys. Double-click to reset.
4. **Space**
   The reverb. Drag up for more of it, and sideways for a longer tail.
5. **Level**
   The loudness of the whole instrument. Drag up for louder.
6. **Pulse**
   A held note plays itself again in a rhythm, locked to the tempo of your song. Drag up for more (0 is off), sideways for the mix between the held notes and the rhythm. Click it to open the pulse menu.
7. **Layer**
   Something that rings on top of the wave when a note starts. Tap to change it (none, tine, bell, glass, wood), drag up for more of it.
8. **Tune**
   Puts every tone of the sound you dropped in onto a note of the scale, so it stays in tune with what you play. From halfway up it is exact; below that it eases off.
9. **Harmony**
   Opens the harmony menu: the synth and the world each get their own intervals above every key you play. Drag up or down for Purity, which takes the world's own harmony out until only its main pitch is left.
10. **Tilt**
    A seesaw between the synth and the world. Lean it to one side and that one is louder while the other is quieter. Double-click to level it.
11. **The wave**
    The wave under everything: sine, triangle, saw or square.
12. **Keys**
    Splits the keyboard into zones, each played by the synth, the world, or both.
13. **World**
    Click to open the world browser, or drop a sound, a recording or an ambience here (or anywhere on the window). The x takes it out again, back to the quiet drone.

Operation

step by step

1. Put the plug-in on an instrument track and play some notes. You will hear the synth and see the ripples.
2. Bring in a World: click the ring at the top right and choose a sound or a recording in the browser, or drop a file anywhere on the window. The keys now play it in tune.
3. Drag the World drawing sideways to stretch it into a soft cloud, and up for more of

Controls and settings

range · start

Control	Range	Starts at	What it does
World	drag up, sideways	halfway	How much of your sound is heard, and (sideways) how far it is stretched, from 1 to 1000 times slower.
Space	drag up, sideways	halfway	Reverb amount, and the length of its tail.
Level	drag up	–6 dB	The loudness of the whole instrument.
Pulse	off to full; mix	halfway	How much a held note repeats itself in a rhythm; sideways, the mix between the held notes and the rhythm.
Pulse menu	rhythm, scale, rate, style		Sixteen rhythms on two pages (Drifting and Grooves), the scale for the harmony, the rate (locked to your song) and the style.
Rhythm by note	on / off	off	Each of eight zones (the seven white notes, and the black notes together) plays its own rhythm, the same in every octave.
Harmony	a switch for each interval	none	Intervals above every key (the 8ve, 5th, 3rd, 6th, 2nd and 7th), one set for the synth and one for the world.
Purity	drag up or down	off	Takes the world's own harmony out until only its main pitch is left.
Tune	drag up or down	halfway	Puts every tone of the world onto a note of the scale.
Layer	none, tine, bell, glass, wood; amount		What rings on top of the wave when a note starts.
Wave	sine, triangle, saw, square		The wave under everything.
Tilt	synth to world	level	A seesaw: one louder, the other quieter.
Keys	up to four zones	one zone	Splits the keyboard into zones

How it works

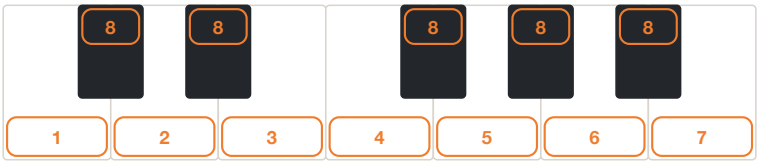
a closer look

KA Chladni has two voices that play together: a synth (a wave, with something that rings on top of it) and the world, a sound or recording you bring in. The world is not played faster as you go up the keyboard. Its pitch is worked out, and every note moves it to that note at its own speed, so a recording of a piano stays a piano.

Stretch plays the world as an endless soft cloud: the sound is rebuilt from its own spectrum, up to a thousand times slower, and still follows the keys. Tune puts the tones of the world onto notes of the scale, so a recording that is a little off the pitch comes into tune with what you play.

Pulse plays a held note again in a rhythm locked to your song's tempo. The harmony adds intervals above every key from a scale. With Rhythm by note, each of the seven white notes and the black notes together make eight zones, and each zone plays its own rhythm, the same in every octave.

RHYTHM BY NOTE: EIGHT ZONES IN EVERY OCTAVE



Zones 1 to 7: one for each white note. Zone 8: all the black notes together.

Each zone can play its own rhythm. The zones are the same in every octave, so the same note always plays the same rhythm, wherever it is on the keyboard.

Or let every note play the same rhythm: switch By note off.

FIG. 47 The octave is cut into zones; each zone can have its own rhythm.

NO. 21 · INSTRUMENT

KA Root

A drum machine that grows grooves from the hits you plant, or from a recording.

KA Root is a drum machine. Plant a few hits on the grid and roots grow from them: answers a few steps away that change from bar to bar, so the beat has a life of its own. Or drop in a recording and Root finds the hits in it and builds a kit from them. It has the classic step grid too, with the steps between the beats, and a stutter or glitch for any step.

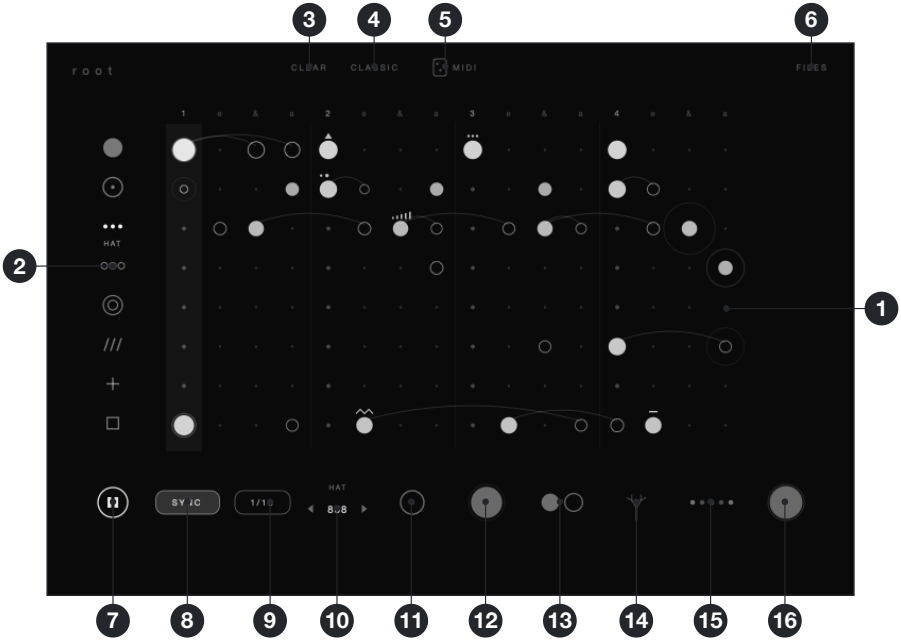


FIG. 48 The parts of KA Root. The numbers are explained on the next page.

Parts

fig. 48

1. **The grid**
   Eight drums, sixteen steps: the beats 1, 2, 3, 4 and the steps between them (e, &, a). Click a step to plant a hit, drag along a row to paint hits, click a hit to take it away, drag a hit up or down for how hard it is. Right-click a step for stutters and glitches. Hollow circles are the roots growing from your hits.
2. **The drums**
   Kick, snare, hat, open hat, tom, clap, rim and perc. Click one to choose it for the controls below. On hover each shows M (mute), S (solo) and an x that clears just that drum's row.
3. **Clear**
   Takes every planted hit off the grid.
4. **Classic**
   Plants a plain, classic beat on the grid to start from (it replaces what is there).
5. **MIDI**
   Drag the beat out as a MIDI file onto a track in your music software: four bars, with the roots as they grow. Root stops playing the pattern itself when you pick it up, so it is not heard twice. A click saves the file on your Desktop.
6. **Files**
   Opens a file browser beside the window. Drop a recording on the window and Root finds the hits in it and builds a kit; drop a single sound on a drum's row to give that drum its own sound.
7. **Play**
   Runs and stops the machine when your song is stopped. While your song plays it turns the pattern off and on.
8. **Sync / Free**
   Click to switch between your song's tempo (Sync) and Root's own (Free). In Free, drag up and down on it to set the tempo.
9. **Rate**
   How fast the steps go: 1/8, 1/16, 1/32, and the triplets 1/8T and 1/16T. Click to step through them.
10. **The sound**
    The sound of the chosen drum: six built-in sounds, and a seventh, your own, from a recording or a file you dropped. Click the arrows to step through them.
11. **Tune and decay**
    Of the chosen drum: drag sideways to tune it up or down, up to make it ring longer.
12. **Tone and level**
    Of the chosen drum: drag sideways from dark to bright, up for louder.
13. **Pan**
    Where the chosen drum sits, left to right.
14. **Roots**
    How much grows from the hits you plant: answers a few steps away that change from bar to bar. Nothing means only what you planted plays.
15. **Swing**
    Drag up to make the off-beat steps late.
16. **Level**
    The level of the whole machine.

Operation

step by step

1. Put the plug-in on an instrument track. Press Play to run it, or play your song with Sync on and it follows your tempo.
2. Click Classic for a plain beat to start from, or click steps on the grid to plant your own hits. Drag along a row to paint hits.
3. Drag the Roots drawing up until the beat starts to answer itself. Down to nothing

Controls and settings

range · start

Control	Range	Starts at	What it does
Roots	0 to 100 %	35 %	How much grows from the hits you plant.
Swing	0 to 100 %	0 %	How late the off-beat steps are.
Level	–24 to +6 dB	0 dB	The level of the whole machine.
Sync / Free	Sync, Free	Sync	Your song's tempo, or Root's own (Free).
Tempo	40 to 220 bpm	100 bpm	Only in Free: drag up and down on the Free pill.
Rate	1/8, 1/16, 1/32, 1/8T, 1/16T	1/16	How fast the steps go.
Sound	six built-in, and yours	the first	The sound of the chosen drum.
Tune	–12 to +12 semitones	0	Of the chosen drum.
Decay	short to long	halfway	How long the chosen drum rings. Only ever

How it works

a closer look

Root plays a grid of eight drums and sixteen steps. What you plant on it plays exactly. Roots are added to it: after a planted hit, answers come a few steps away, different for each drum and different every bar, with a fill on the fourth bar, and a little laid back. They are worked out from the pattern and the bar number, so the same pattern grows the same way every time.

A recording you drop in is cut at the places where a new sound starts. Each piece is measured (how bright, how low, how long) and given to the drum it fits best, so the kit comes from the recording. Sounds you bring are played at their own level and tone until you change them.

Each step can carry an effect, which makes several small hits inside the step, and the hits are put in a list ahead of time so the timing is as exact as the steps.

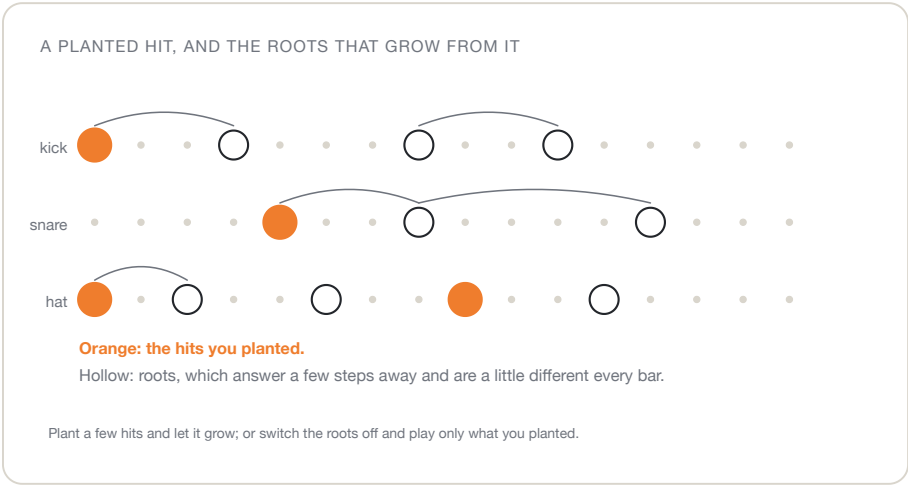


FIG. 49 Roots grow from the hits you plant, a few steps away.

NO. 22

KA Zing

Two knobs, five modes, one button, and zero explanations.

KA Zing has five modes, two knobs and a button. Will we tell you what any of them do? No. Put it on a track, turn Amount up, and use your ears. That is what they are for.

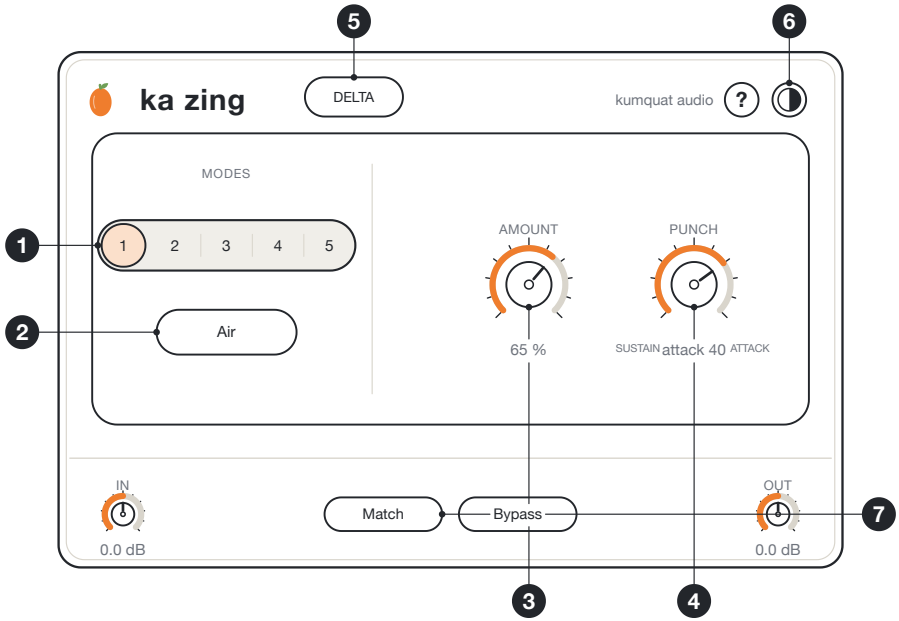


FIG. 50 The parts of KA Zing. The numbers are explained on the next page.

Parts

fig. 50

1. **Modes**
   Five of them. Our lips are sealed.
2. **Air**
   Press it. You will know.
3. **Amount**
   More is more. At 0 you hear nothing new.
4. **Punch**
   Left and right are not the same. Go and find out. The middle keeps its hands to itself.
5. **Delta**
   Plays only what the plug-in adds to the sound.
6. **Help and light/dark**
   The ? turns the pop-up help on and off. The half circle switches light and dark.
7. **In, Out and Match**
   In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

Operation

step by step

1. Put it on a track. Any track. Play something.
2. Pick a mode, 1 to 5. Yes, the buttons are just numbers. Try them all, and keep the one that makes you grin.
3. Turn Amount up from 0. Not feeling it? Try another mode before you blame us. Press Delta if you want to eavesdrop on what it adds, and Match so a bit of extra volume does not fool you.
4. Turn Punch left. Turn it right. They are not the same, and no, we are not saying how. Press Air when you feel fancy.
5. Flip Bypass to check that it is better and not just louder. If it is only louder, back off.

Controls and settings

range · start

Control	Range	Starts at	What it does
Modes	1 to 5	1	Five of them. Our lips are sealed.
Amount	0 to 100 %	0 %	More is more. At 0 you hear nothing new.
Punch	left to right	middle	Left and right are not the same. Go and find out. In the middle it keeps its hands to itself.
Air	on / off	off	Press it. You will know.
Delta	on / off	off	Plays only what the plug-in adds. Nosy, but allowed.
In / Out	–24 to +24 dB	0 dB	The two knobs under the meters: before and after the plug-in.
Match	on / off	off	Keeps the output as loud as the sound that came in, whatever In and the plug-in do. See In, Out and Gain Match.

TIP Both knobs at the start, and Air off, leave the sound exactly as it was. So you can always crawl back.

TIP A little goes a long way. Trust your ears, not your eyes: there is nothing to look at, on purpose.

CAUTION It is a bit addictive. Everything else in the mix will start to sound dull next to it. Check the whole mix once in a while, and take a break.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo
Delay	120 samples (2.5 ms at 48 kHz)
Window	580 × 290, can be resized

NO. 23

KA Contour

Evens out a voice by itself: turns down boomy, nasal and harsh bumps and fills in holes, with a high pass.

KA Contour evens out a voice. It listens for a few seconds, finds the narrow peaks and valleys in the sound of the voice (a boomy note, a nasal honk, a harsh ring, a hole in the middle), and turns the peaks down and fills the valleys in a little. It does this by itself, in the background, with one main knob.

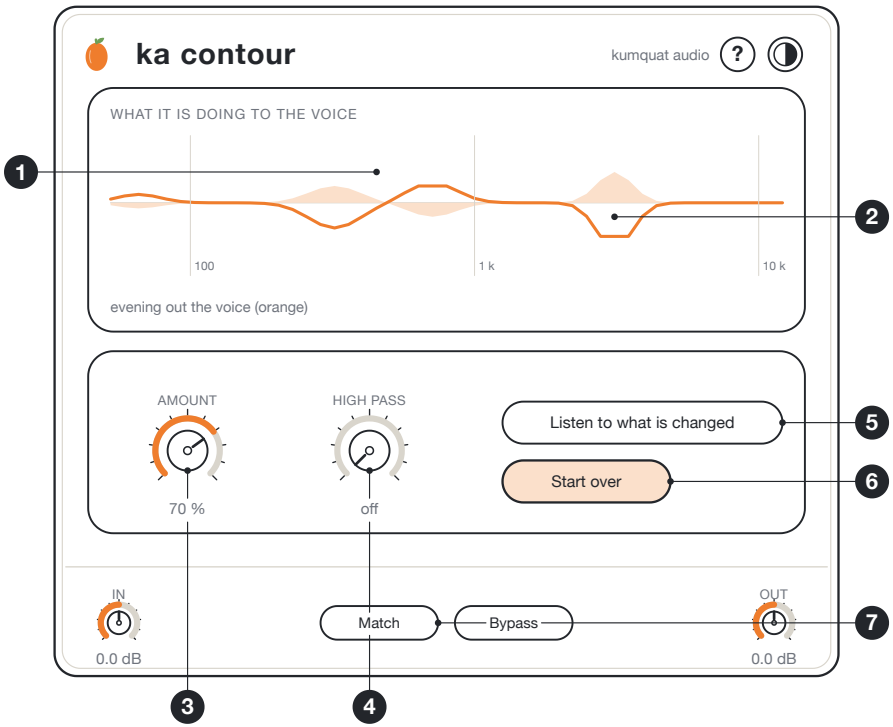


FIG. 51 The parts of KA Contour. The numbers are explained on the next page.

Parts

fig. 51

1. **What it found**
   The faint shape shows how far each part of the voice sits above or below its own smooth line.
2. **What it does**
   The orange line is the correction: where the voice sticks up it is turned down, where there is a hole it is turned up. It never moves faster than 1.5 dB a second.
3. **Amount**
   How closely the voice follows its own line.
4. **High pass**
   Takes out rumble below the corner, with a gentle slope. Off at the far left.
5. **Listen**
   Plays only what is being changed, and what the high pass takes out.
6. **Start over**
   Forgets the voice and learns it again.
7. **In, Out and Match**
   In turns the sound up or down before the plug-in hears it, and Out turns the result up or down. Bypass plays the sound without the plug-in. Match keeps the output as loud as the sound that came in, whatever In and the plug-in do, so you hear what a setting does and not just a change in loudness.

Operation

step by step

1. Put the plug-in on one voice, a dialogue or a vocal, and play it. The picture says it is listening and counts up to 100 %. That takes about three seconds of talking or singing.
2. Watch the picture. The faint shape shows how far each part of the voice sits above or below its own smooth line. The orange line is what Contour is doing about it: where the voice sticks up it is turned down, where there is a hole it is turned up.
3. Turn Amount to taste. At 70 % the voice follows its line closely without changing its character. Press Listen to hear only what is changed: it should sound like a faint, thin version of the voice, and never loud.
4. Turn High pass up to take out rumble and thumps: a hand on the stand, traffic, an air conditioner. Start at about 80 Hz. Listen plays what it takes out as well.
5. When the voice on the track changes, press Start over. Otherwise Contour follows the change by itself over about ten seconds of the new voice.

Controls and settings

range · start

Control	Range	Starts at	What it does
Amount	0 to 100 %	70 %	How closely the voice follows its own smooth line. 0 leaves it as it is.
High pass	20 to 400 Hz	off	Takes out rumble below this frequency with a slope of 18 dB an octave, and adds no delay. Off at the far left.
Listen	on / off	off	Plays only what is being changed, and what the high pass takes out.
Start over	button		Forgets the voice and learns it again.
In / Out	In ±24, Out ±12 dB	0 dB	The two knobs under the meters: before and after the plug-in.
Match	on / off	off	Keeps the output as loud as the sound that came in, whatever In and the plug-in do. See In, Out and Gain Match.

TIP Give it a few seconds of the voice before you judge it. While the picture says it is learning, the correction is still growing.

CAUTION Contour is made for one voice at a time. Two people on one track are averaged together.

CAUTION The plug-in delays the sound by 85 ms. Your music software lines everything up. Do not compensate by hand.

Specifications

Formats	Audio Unit, VST3
Channels	mono or stereo
Delay	4096 samples (85 ms at 48 kHz)
Window	560 × 400, can be resized

How it works

a closer look

KA Contour measures how loud every part of the spectrum is while the person is speaking or singing. It does this in 46 bands, a sixth of an octave apart, from 70 Hz to 12.7 kHz. Pauses, breaths and hiss are left out: only the parts that are clearly voice count.

From that average it draws a smooth line that ignores anything narrower than about an octave. A boomy note, a nasal honk or a harsh ring sticks up through the line, and a hole in the middle falls under it. Contour turns the peaks down, by up to 8 dB, and fills the valleys in, by up to 4 dB. A dark voice stays dark and a bright voice stays bright: only the bumps are evened out.

It starts after about a second of speech and is fully working after about three. The correction then moves by at most 1.5 dB a second, so nothing pumps, and it stays where it is while nobody speaks. The average covers about the last twelve seconds of speech, and it is saved with your project.

The high pass is separate from all this. It takes out what is below its corner, with a gentle slope of 18 dB an octave, and does not delay the sound.

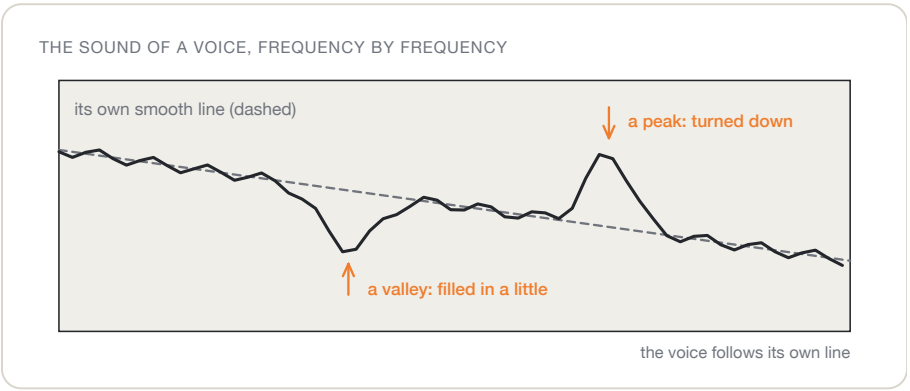


FIG. 52 A narrow peak is turned down and a valley is filled in, so the voice follows its own smooth line.

ARA

Plug-ins that read whole recordings

KA Silence Trim, KA Drum Gate, KA Auto Mixer and KA Time Align read whole recordings, which is something only a few music programs allow: it is called ARA. They come as Audio Units only.

Today they need Logic Pro to be opened with Rosetta. Select Logic Pro in the Finder, press Cmd+I, tick “Open using Rosetta”, and open Logic again. The other plug-ins do not need it.

Put each one first in the Audio FX list of its track. Select a region and press play: the window fills as the plug-in reads the audio.

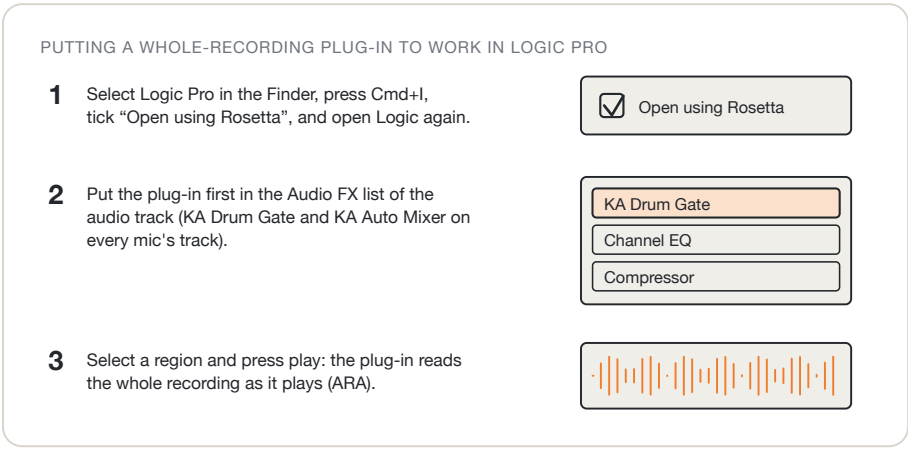


FIG. 53 Three steps to put one of them to work.

CAUTION If the window says it is not connected to Logic’s audio, one of the three steps has been missed: Rosetta, the first place in the list, or a region that is selected and playing.

CARE

Care and troubleshooting

If this happens	Do this
A plug-in does not appear in the list	Restart the music software. In Logic Pro, open the Plug-in Manager, select the Kumquat plug-ins and choose Reset & Rescan Selection.
An old name is still showing	Quit Logic with Cmd+Q and open it again. If it persists, remove the old copies from /Library/Audio/Plug-Ins/Components.
My key was not accepted	Copy all of it, from KA1 to the end, with no spaces added or lost. A key for one plug-in does not unlock the others.
The pop-up help does not show	Press the “?” at the top of the window to turn it on. It is one choice for all the plug-ins.
Which part of the picture does a knob change?	Point at the knob, or hold it down while you turn it: the part of the picture it controls lights up, with the knob’s name and value.
KA Drum Gate, Silence Trim, Auto Mixer or Time Align shows a message instead of a picture	Open Logic Pro using Rosetta, put the plug-in first in the Audio FX list, select a region and press play.
The sound is late	KA De-Click, De-Verb, De-Noise, Contour, Mic Blend, Saturator (it oversamples), EQ (with the suppressor), Multiband (in Linear phase) and Limiter delay the sound a little and tell the music software, which lines everything up. Do not compensate by hand. For playing through a plug-in, the Saturator, EQ and Multiband have a Live switch that takes the delay away; the compressors, De-Esser, Delay, Reverb and Rider have no delay at all.
A window is too small or too big	Drag the corner at the bottom right. The size is remembered, and it is the same for all the plug-ins: make one bigger and the others open that big.
The Limiter pumps	Leave Release on Auto, or set 60 to 100 ms. A release of a millisecond or two is the dirtiest setting.

If something is not right, write to **audiokumquat@gmail.com** and say which plug-in, which music software, and what you were doing.

REFERENCE

Specifications at a glance

Plug-in	Formats	Delay at 48 kHz	Window
KA De-Esser	AU, VST3	8 ms (none with Live)	520 × 560
KA De-Click	AU, VST3	1152 samples · 24 ms	520 × 406
KA De-Noise	AU, VST3	2048 samples · 43 ms	700 × 520
KA Contour	AU, VST3	4096 samples · 85 ms	560 × 400
KA De-Verb	AU, VST3	2048 samples · 43 ms	520 × 406
KA EQ	AU, VST3	none · 85 ms with the suppressor	1000 × 726
KA Limiter	AU, VST3	746 samples · 15.5 ms	860 × 688
KA Mic Blend	AU, VST3	484 samples · 10 ms	520 × 456
KA Saturator	AU, VST3	61 samples · 1.3 ms	560 × 496
KA Delay	AU, VST3	none	600 × 496
KA Bus Comp	AU, VST3	none	660 × 524
KA Reverb	AU, VST3	none	860 × 542
KA Zing	AU, VST3	120 samples · 2.5 ms	640 × 292
KA Drum Gate	AU (ARA)	reads whole recordings	820 × 720
KA Silence Trim	AU (ARA)	reads whole recordings	520 × 548
KA Auto Mixer	AU (ARA)	reads whole recordings	760 × 720
KA Time Align	AU (ARA)	reads whole recordings	920 × 760

They all run on macOS on Apple silicon and Intel, as universal builds, and are signed and notarized by Apple. They work with mono and stereo tracks. Window sizes are in points; every window can be resized and keeps its shape.

Delay: the plug-ins that look ahead tell your music software how late they are, and it lines everything up. KA Limiter's delay stays the same whatever Look-ahead is set to.

REFERENCE

Glossary

Ceiling
The highest level a limiter lets through.

Clipper
Shaves off the very tops of peaks at once, instead of turning the whole sound down.

De-ess
To turn down the harsh part of “s” and “sh”.

dB
Decibel: a way of measuring loudness. 6 dB is twice or half the height of a wave.

Bleed
One microphone hearing something meant for another.

ARA
A way for a plug-in to read a whole recording instead of only what is playing.

LUFS
Loudness as it is heard: momentary is the last 0.4 s, short-term 3 s, integrated everything so far.

Look-ahead
How far ahead a limiter looks, so that it can turn down before a peak arrives. It delays the sound by as much.

Polarity
Whether a wave goes up first or down first. Two mics can be upside down to each other.

All-pass filter
Changes only the phase of a sound, not its loudness. KA Time Align uses one to line up a mic’s low end.

Q
How narrow an EQ band is. Higher is narrower.

Gain matching
Keeping the output as loud as the input, so that a setting is judged by how it sounds and not by how loud it is.

Harmonics
Extra tones above a sound, at two, three, four times its pitch and so on. Saturation makes them.

Oversampling
Working at a higher sample rate for a moment, so that new high frequencies do not fold back as noise.

Ratio
How firmly a compressor turns the sound down once it is over the threshold. At 4:1, 8 dB over comes out 2 dB over.

Feedback
Sending what comes out of an effect back into it, so that it repeats.

Pre-delay
The wait between a sound and the start of its reverb.

RT60
The time a reverb takes to fall by 60 dB. It is what Decay sets.

Crossover
The frequency where one band of a multiband plug-in ends and the next begins.

Linear phase
A way of splitting the sound into bands without turning its phase, at the cost of a delay.

True peak
The highest level the sound reaches between its samples once it is played. Streaming services measure it.

Sibilance
The hissing sound of “s”, “sh” and “t”.

Sidechain
A second input that tells a plug-in what to listen to.

True peak
The height of the wave between the samples,

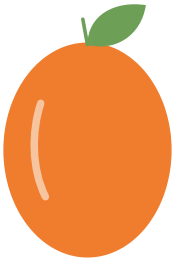
REFERENCE

Index

Aliasing 54
Align 38
Align all 70
All (FET ratio) 79
All-pass filter 70
ARA 106
Attack 32
Auto (de-esser) 8
Auto release 62
Band (multiband) 75
Band edges 8
Bass 32
Bite 37
Bleed 42
Ceiling 32
Clean (mode) 32
Click counter 12
Clip 32
Color (compressors) 79
Cross 58
Crossover 75
Decay 66
Delay to add 70
Depth 42
Drive 54
Dynamic band 26
Early reflections 66
Emphasis 83
Expand 75
Export 46
Export aligned 70
Fade time 50
Feedback 58
Focus 21
Linear phase 75
Listen 8
Look-ahead 32
Loud (mode) 32
Make room 30
Makeup 62
Minimum phase 75
Mix (compressor) 62
Over music 87
Peak reduction 83
Ping-pong 58
Pre-delay 66
Pre-roll 46
Pulse 91
Punch (mode) 32
Range 8
Range 75
Ratio 62
Release 32
Riding 87
Roots (drums) 95
Rosetta 106
RT60 66
Sensitivity 12
Size 66
Soft, Warm, Hard 54
Space (reverb) 66
Split bass 36
Step effects 95
Stereo (delay) 58
Stretch 91
Suppressor 30
Synced time 58
Threshold 32

FOR YOUR OWN NOTES

Notes



Kumquat Audio

kumquataudio.com

Audio Unit is a trademark of Apple Inc. VST is a trademark of Steinberg Media Technologies GmbH. Logic Pro and macOS are trademarks of Apple Inc. All other names belong to their owners.

Edition 1.7.11 · October 2026