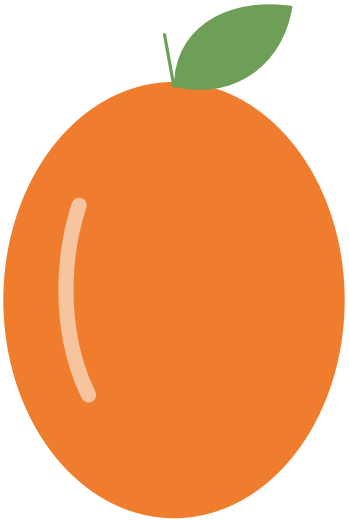


KA De-Noise

Instruction manual

Kumquat Audio · Edition 1.7.11 · October 2026

kumquataudio.com

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.

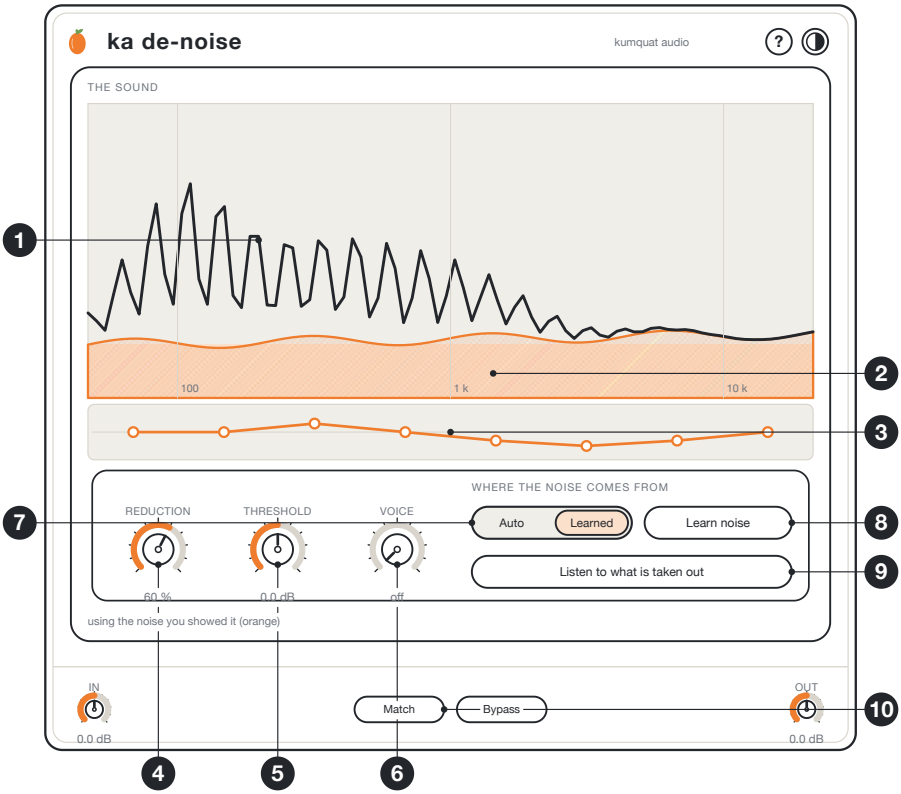


FIG. 1 The parts of KA De-Noise. The numbers are explained on the next page.

Parts

fig. 1

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.

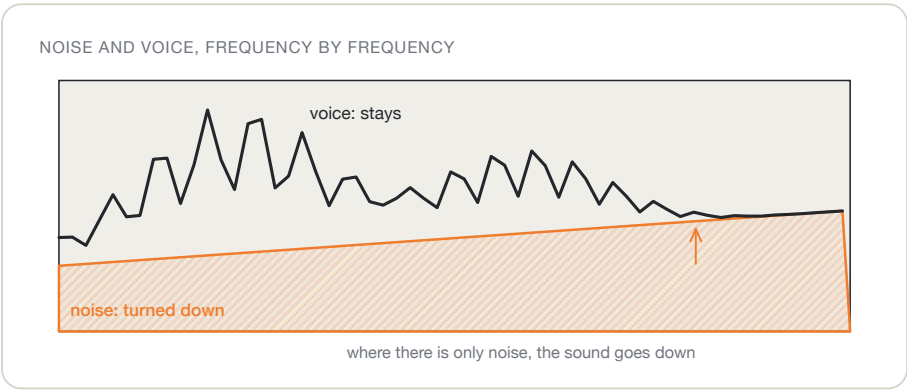


FIG. 2 Where there is only noise, the sound is turned down.

Credits and notices

what is inside

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.

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.

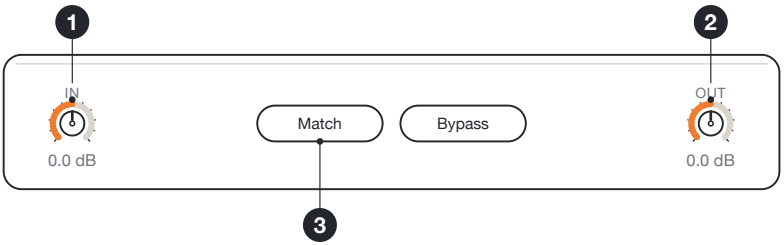


FIG. 3 The strip along the bottom of a window.

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.

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.

FOR YOUR OWN NOTES

Notes