manual

neutone
morpho

Plugin ver. 1.0.0
Last updated on 8th April 2024

https://neutone.ai/morpho

contents

overview

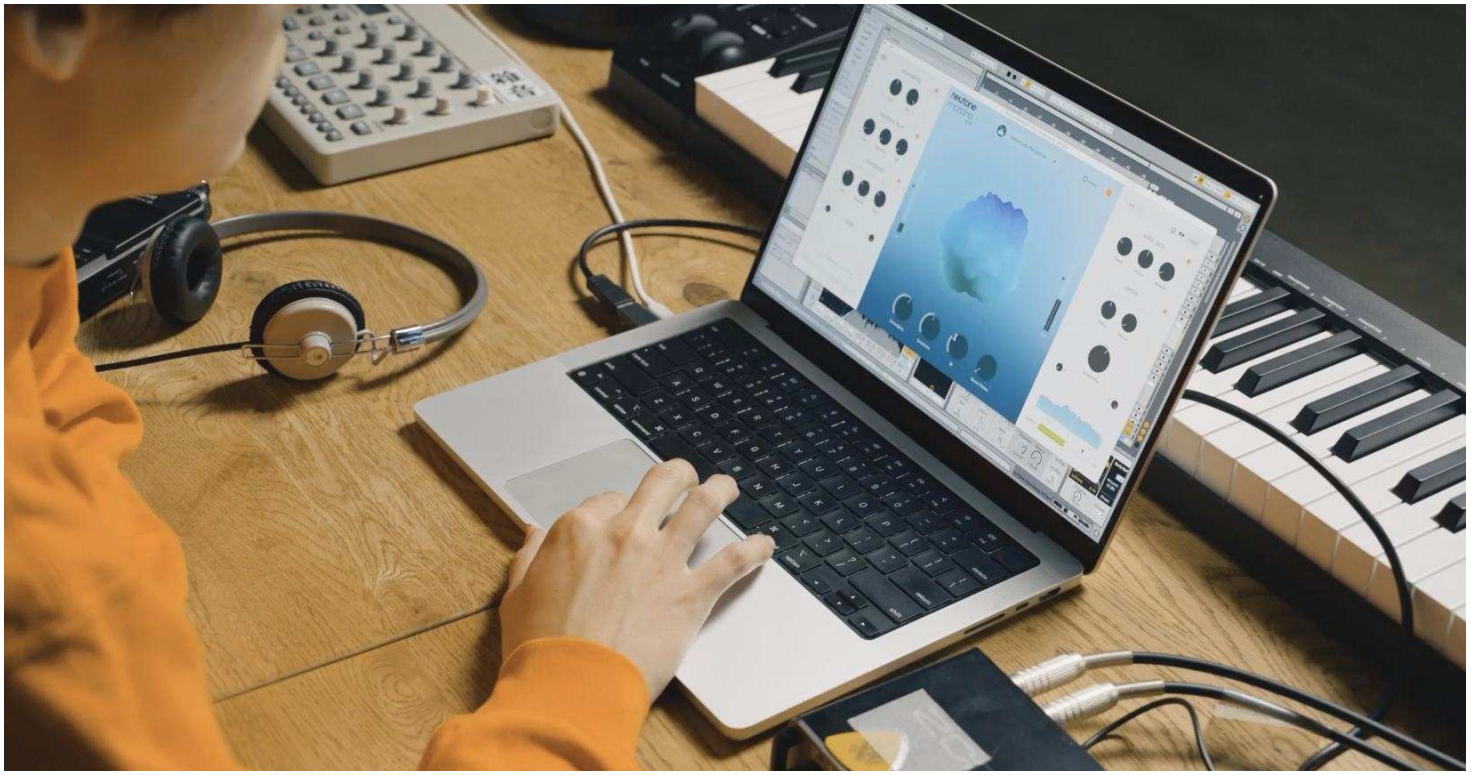

We're pleased to present Neutone Morpho, a realtime tone morphing plugin. Our cutting-edge machine learning technology can transform any sound into something new and inspiring.

Neutone Morpho processes audio directly, bypassing the traditional MIDI convention based on pitch and velocity. This approach enables the capture of even the most subtle nuances from your input audio.

Additional models are also available for purchase in our model store, accessible directly in the plugin. For those interested in creating their own Morpho model, we will soon be launching Neutone Cocoon, our training service where you can upload your audio materials.

We're excited to see what you'll create with Neutone Morpho and Neutone Cocoon*!* Let's begin exploring the limitless creative potential that Neutone offers*!*

specifications /installation

Mac

MacOS 10.15+ for Intel
Mac OS 12.0+ for M series
Plugin Format: VST3/AU

Windows

Windows 10 & Windows 11
Plugin Format: VST3

Tested DAWs

Ableton Live 11 & 12
Logic Pro 10
Bitwig Studio 5
Studio One 6
Cubase 13

Installation

1. Download the installer for your platform
2. Run the installer

(you need at least 8GB of RAM / 1GB of storage)

💡 If you're using an older machine, consider reducing the resolution of the visualizer or turning it off in the settings to conserve CPU power.
Also make sure to use a larger buffer size, especially when exporting or bouncing audio through the plugin it prevents glitchy or noisy output.

plugin overview

Main Menu
↳Model Browser (p.9)
↳Information
↳Settings

Orange Button
enables/disables processing feature

Quick Model Selector
to switch between downloaded models

Presets
choose or save presets per model

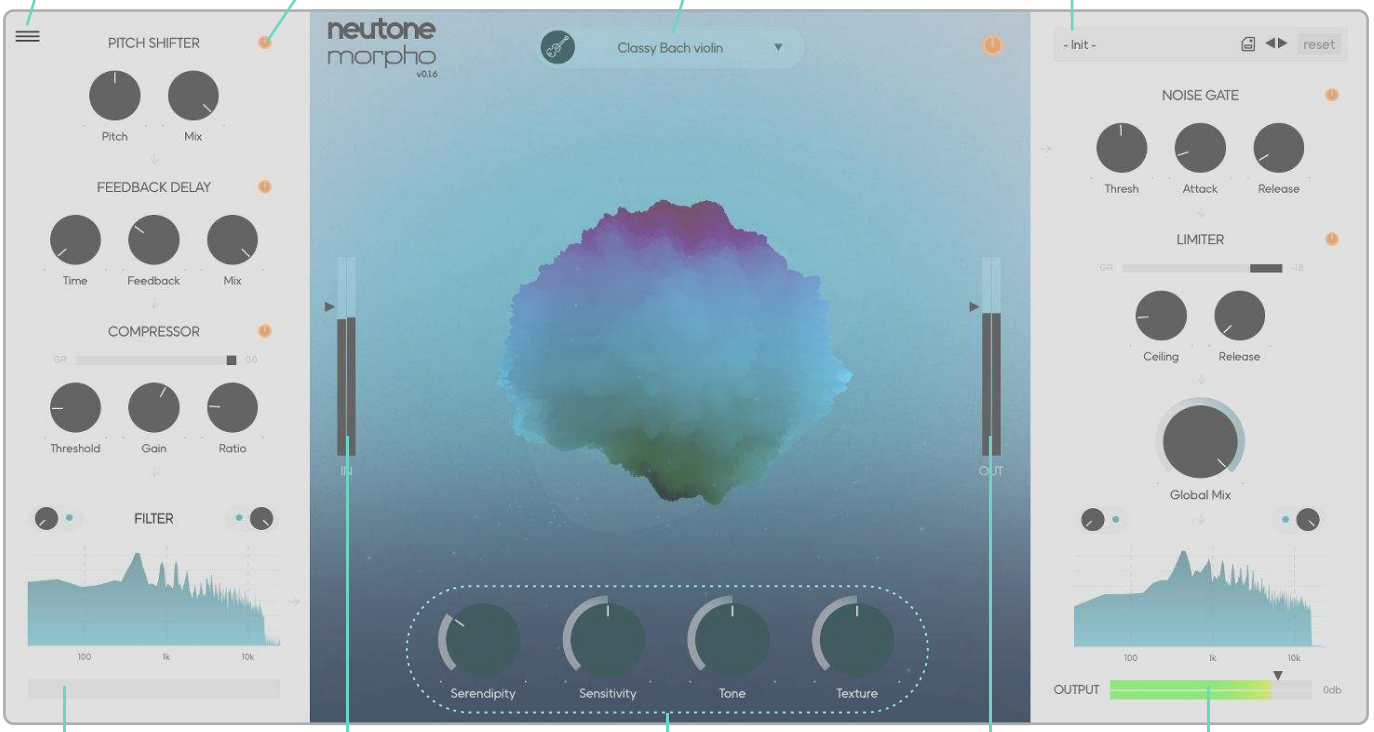

Pre-processing (p.10)

Post-processing (p.11)

Info Box
Hover over elements to see more information on them here

Model Input Gain
Gain for the model (changing it can also change the timbre not only it's loudness)

Model Parameters (p.6)
Macro knobs to change the model behaviour (click the central orb to switch into micro mode)

Model Output Gain
output of the morpho model

Main Output
global output signal

About

In the upcoming pages we will go through the interface and explain each of the elements. Generally Neutone Morpho is split into 3 sections: pre-processing on the left side, which is then processed by the model in the center and a post-processing section to shape the output on the right.

model parameters

Morpho Model

At the core of Neutone Morpho are the Morpho AI models, where the magic happens. You can interact with a loaded Morpho model in two modes to influence the tone-morphing process.

💡 You can use the enable/disable button to toggle the model processing. This allows you to compare the original input sound with the model's output.

Macro View

In *Macro View*, you have four macro knobs that allow you to change the general behaviors of the model. For example, the "Serendipity" knob determines the level of randomness in the transfer, with higher serendipity resulting in more complex sounds. Each model has its own unique macros. For instance, a percussion model may have a knob called "decay" that controls the decay of each hit.

Micro View (Experimental)

The Morpho AI model crunches incoming audio into 6 or more values, referred to as "latent variables." In Micro View, you can directly offset and scale those values with Morpho Modifiers to steer the output in certain directions. You may even uncover a hidden correlation between a latent variable and the pitch of the output! (More technical details are available on p.8.)

💡 When you have larger scale values, your model tends to behave more randomly. If you reduce them, you'll get a more predictable and controlled sound. You can observe the effect clearly by adjusting the Serendipity knob in the macro view.

💡Macro knobs control one or more Morpho Modifiers simultaneously, hence the name. We plan to provide a way for you to define your own macros in the future.

neutone morpho algorithm

Under the hood, Morpho Model utilizes an AI algorithm called “autoencoder” to transform the incoming audio. It basically consists of two networks, the encoder and the decoder. The encoder analyses the incoming audio, outputting information about the pitch, loudness and other more subtle features of the incoming audio (this is sometimes referred to as the latent variable). The decoder then uses this information to figure out how it should synthesize the sound. The actual Morpho model has some extra details and modifications to improve the fidelity and be robust to different input sounds, but that is the big picture.

When creating a Morpho model, the autoencoder is trained on only a certain type of sound. For example, a violin Morpho model is trained on violin performances. This violin model then interprets any sound you input in the context of violins and outputs a violin-like sound, creating a “timbre transfer” effect. If you input your flute performance into this violin model, it will analyze how the flute was played, and reimagine what it would be like if you had been playing a violin instead. This timbre transfer effect can lead to unpredictable and fun outcome if you feed it inputs that are drastically different from the data the model was trained on.

💡 If you’re using an older machine, consider reducing the resolution of the visualiser or turning it off on the settings view to conserve CPU power. And using bigger buffer size helps on older machines.

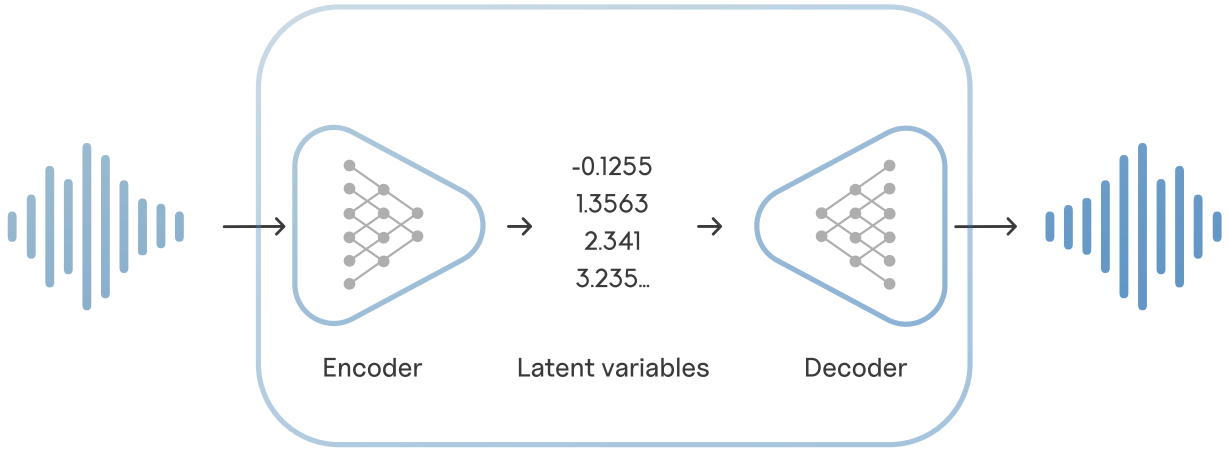

A diagram of neutone morpho algorithm

micro view

💡This page provides an explanation of how the Micro View and Morpho Modifiers function for advanced users. Feel free to skip this page if you wish.

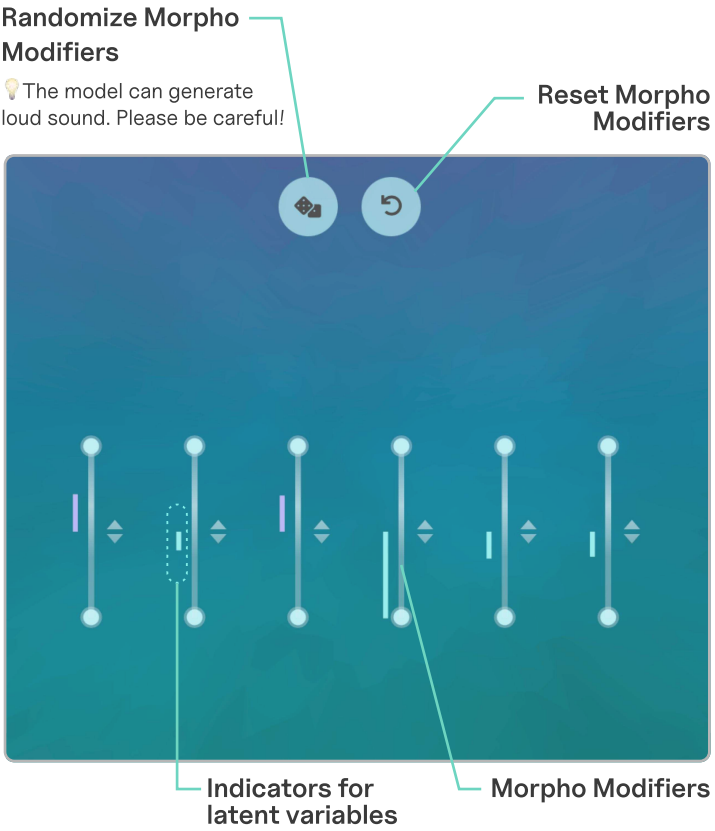

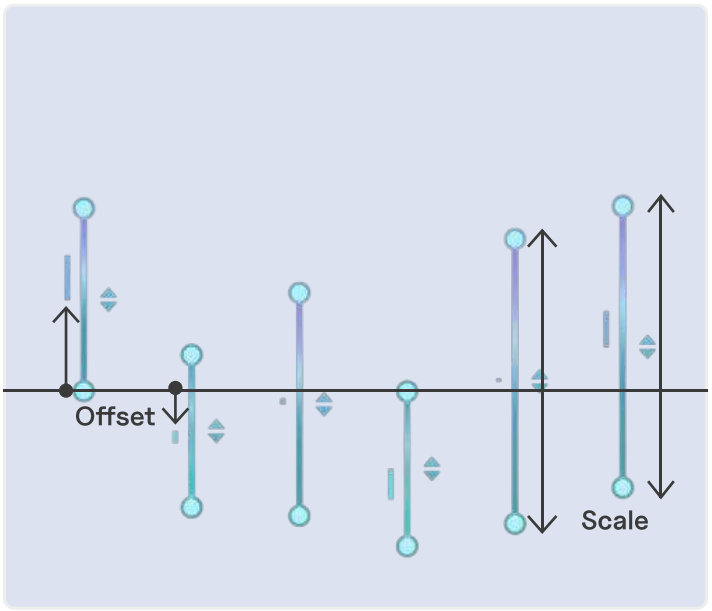

As outlined earlier, Morpho models are trained to condense incoming audio into N (=6 or more) numbers, known as latent variables. These variables are designed to represent the N most significant perceptual characteristics of the sound used during training. Identifying which number corresponds to which characteristic can be challenging, as the machine learning process determines them automatically.

💡Don't worry! We've identified some of the most useful relationships between latent variables and perceptual characteristics of sound and have assigned them to the Macro knobs.

You can use Morpho Modifiers to directly offset and scale each of those values, steering the output in certain directions. The rule of thumb is that the closer latent variables are to zero (vertical center), the more common the output sounds. In other words, you can create extreme sounds using larger offsets and scales. So, keep an eye on latent variable indicators and experiment with Morpho Modifiers to uncover unique and captivating textures within the Morpho model!

model browser and model purchase

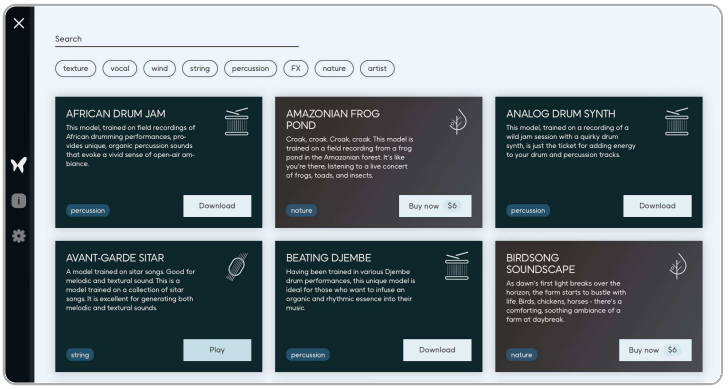

In the main model browser, you can view all models available for purchase, download, and use. You can open the model browser from the side bar and also filter models by categories or tags.

When you open the browser for the first time, you'll be prompted to create a Neutone account (if you haven't already) and sign in with your credentials.

Before purchasing any models, you must buy the full Neutone Morpho version.

Model Purchase

Once you've linked your credit card information to your Neutone account, you can begin purchasing models via the plugin.

If the model doesn't meet your expectations or differs from what you expected, you can request a refund within **three days** of your purchase. Please contact us at support@neutone.ai.

pre-processing

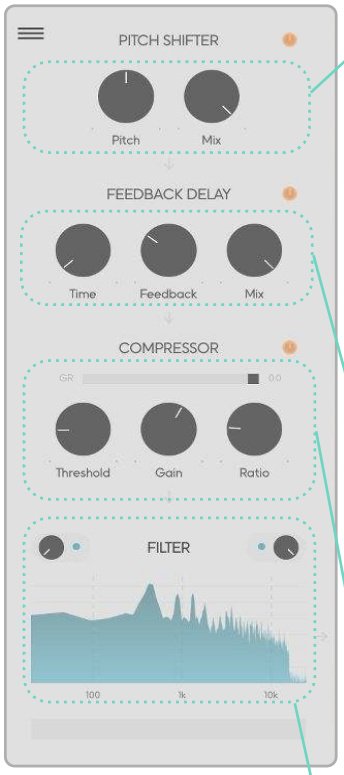

Pitch Shift

You can adjust the pitch of the input sound. By using the mix knob, you can blend the sound at its original pitch with the shifted pitch to create a harmonious effect.

💡 You can explore different sound characters in a Morpho model by the pitch shift. With a choir model, if you increase the pitch you'll start to hear more female voices for example.

Feedback Delay

You can use Feedback Delay with Morpho Model to create a more complex sound with rhythmic variations.

Compressor

To achieve optimal performance from a Morpho model, it is important to stabilize the loudness of the input sound. The compressor feature and input gain meter can assist you in doing so.

💡 Experiment with adjusting the input gain to discover different characteristics in a Morpho model.

Input Filters

Each model has a specific range of input frequencies where it performs optimally. Inputs outside of this range can result in unwanted effects on the output. By using a combination of filters and pitch shifting, you can adjust and restrict the pitch of the input sound to fit within that desired range.

💡 Experiment with adjusting the filter cutoff frequency to discover different characteristics of the model.

post-processing

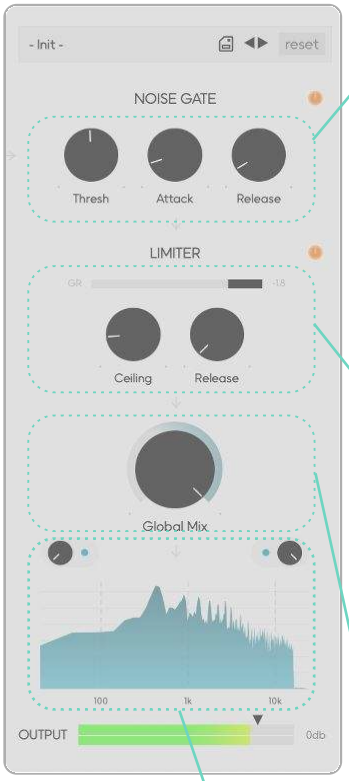

Noise Gate

The noise gate is useful to suppress unwanted tail sound. You can also create staccato effects.

Limiter

Due to the stochastic nature of AI models, you need to ensure that the output sound is at the desired loudness. The limiter to the output is handy for the task.

Global Mix

You can mix “dry sound”(input) and “wet sound” (morphed sound) with this knob.

💡Try blending the wet sound a little bit to add a “flavor” of the trained sound to the input.

Output Filters

Before output, make sure to filter out any unwanted frequency bandwidth and use the most interesting part of the modified sound!

signal flow

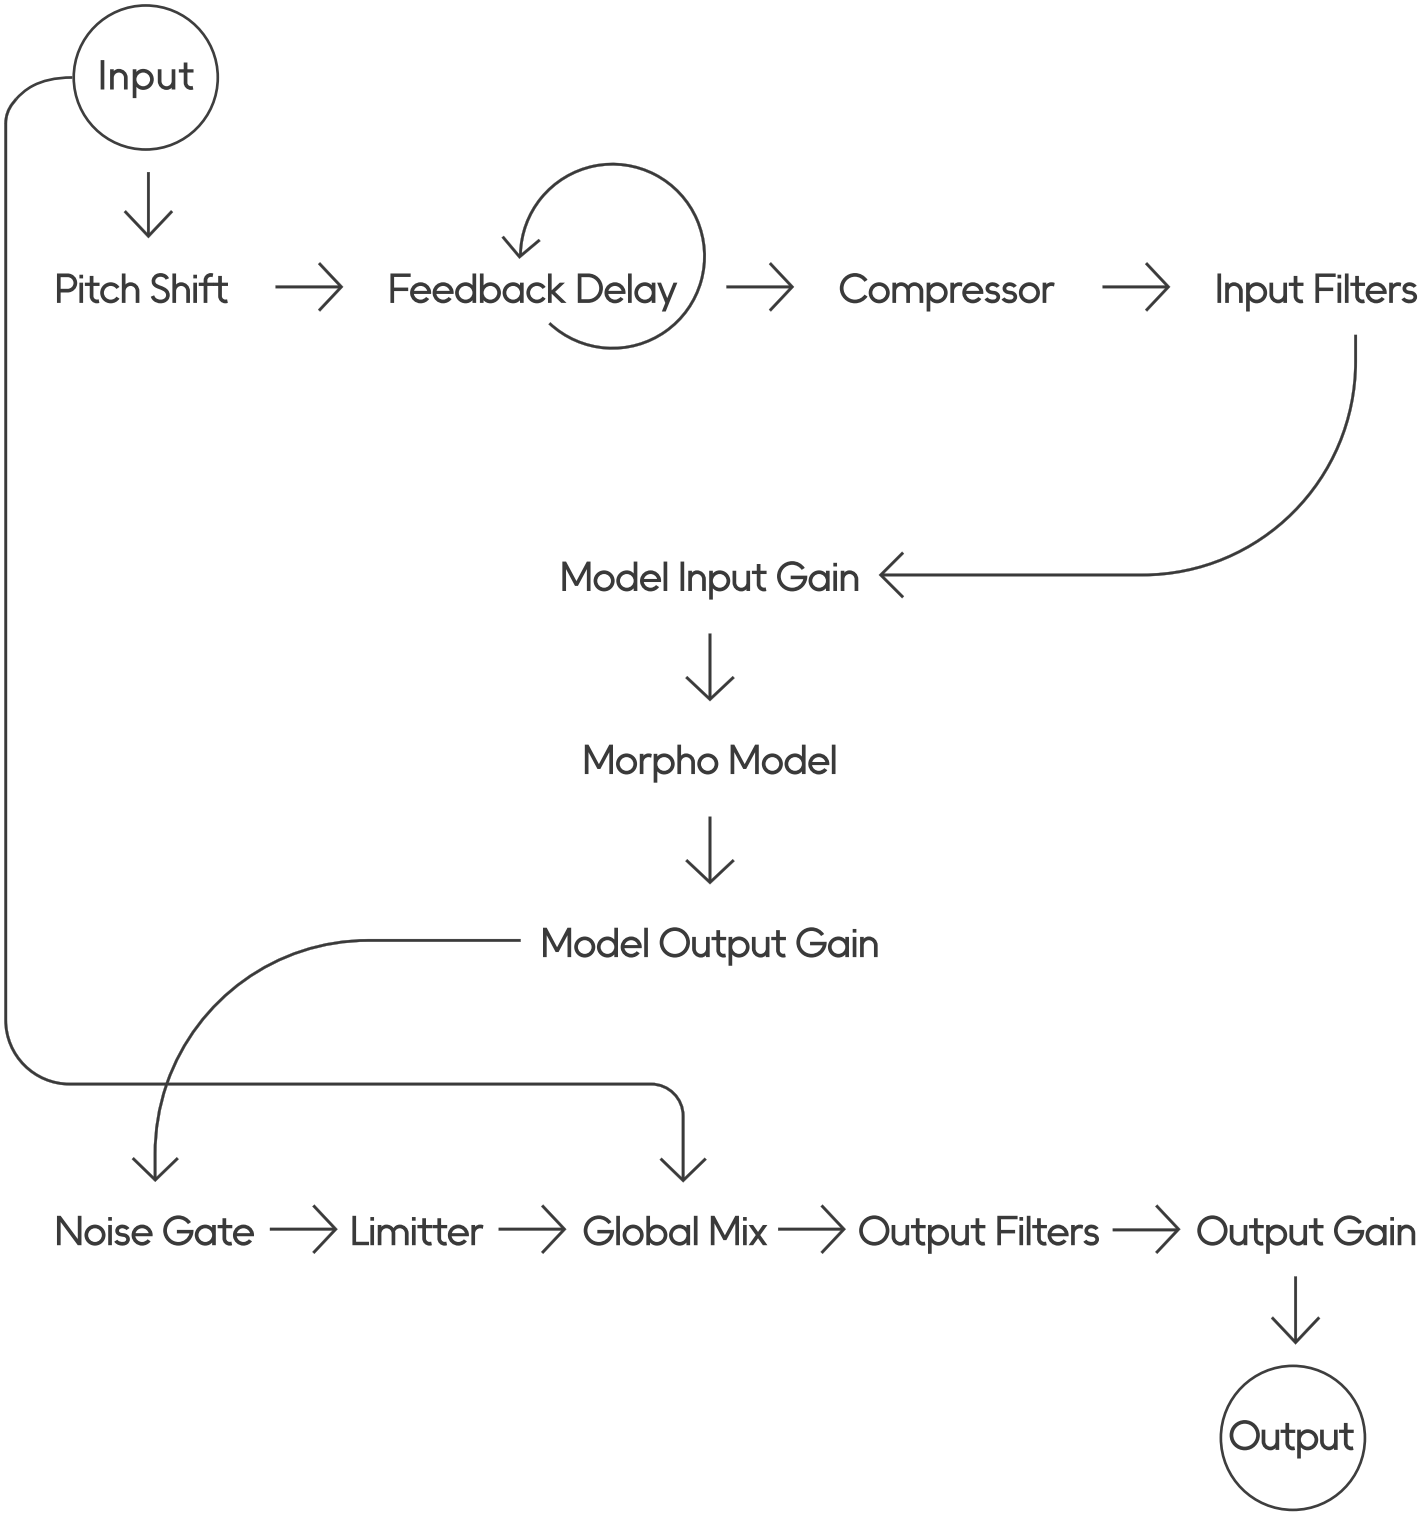
Input
Pitch Shift
Feedback Delay
Compressor
Input Filters
Model Input Gain
Morpho Model
Model Output Gain
Noise Gate
Limitter
Global Mix
Output Filters
Output Gain
Output

ethics

We respect the copyrights of artists and musicians, as well as the efforts they put into creating music and sound materials.

We do not use any copyrighted materials for model training.

Morpho models are trained using publicly available training data that is either in the public domain or licensed under Creative Commons.

For more details, please refer to the model card of each specific model.

Credits

Team Lead	Andrew Fyfe
Plugin Development	Andrew Fyfe Alfie Bradic Christopher Mitcheltree
Frontend Engineer	Nico Pellerin Matthias Schäfer
Backend Engineer	Bogdan Teleaga
Model Development	Naotake Masuda
Model Training	Kay Nohara Francesco Papaleo
Graphic Design	Naoki Ise Matthias Schäfer Mike Sekine
PR	Justin Ruiz
Product Strategy	Stanley Gabriel
Advisor	Kaoru Sugano Hitoshi Miyata
Product Direction	Nao Tokui

Contact /Support

neutone

Visit our website!
https://neutone.ai/

If you have any further questions,
please feel free to contact us at
support@neutone.ai

We have also a dedicated Discord channel
https://discord.com/invite/r6WwYCvJTS

neutone morpho