

contents

overview	3
specifications / installation	4
plugin overview	5
model parameters	6
neutone morpho algorithm	7
micro view	8
model browser and model purchase	9
pre-processing	10
post-processing	11
signal flow	12
ethics	13
credits	14

overview




Thank you for downloading Neutone Morpho, our realtime tone morphing plugin that brings the power of machine learning to your DAW.

Sounds you send into Morpho are resynthesized into a totally different style whilst keeping the overall shape of the audio intact. With this in mind, you can think of Morpho as a synthesizer that is stimulated with sound rather than traditional MIDI information.

To get started, select a Morpho model from the plugin's browser. You will find a variety of pre-trained models included, with additional models available for purchase in the Store. Want something completely unique and personal to you? Train your own custom models with our online model training service.


We're thrilled to have you join us on this new frontier of sound design and sonic exploration. Happy morphing!

specifications /installation



💡 If you're using an older machine, consider reducing the resolution of the visualizer or turning it off in the settings menu. Increasing your DAW buffer size can also improve performance.

plugin overview



About

In the upcoming pages we will go through the interface and explain each of the elements. 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 tone morphing 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 tone morphing 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 utilizes a neural audio algorithm called an “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 outcomes 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 visualizer or turning it off in the settings menu. Increasing your DAW buffer size can also improve performance.



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.

Randomize Morpho Modifiers

💡 The model can generate loud sound. Please be careful!

Reset Morpho Modifiers




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 [my models] tab you can view all models available to you for download. You can filter models by categories, tags or favorites.

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 need to own the full edition of Neutone Morpho.



Model Purchase


The [store] tab allows you to view and purchase additional models.

Once you've linked your credit card information to your Neutone account, you can purchase 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.jp

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 characteristics of the Morpho model via the pitch shifter. For example, increasing the pitch when using a choir model might produce more female voices.

Feedback Delay

You can use the Feedback Delay module to create more complex sounds 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 sounds. You can also use it to create staccato effects.

Limiter

The volume of the model's output can vary greatly depending on the settings. The limiter can help prevent large volume spikes and clipping.

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



ethics

We respect the copyrights of artists and musicians, and train models exclusively on sounds for which we have the rights to use. We cite all our sources and display them in the plugin's model browser.

[Read more about our stance on AI and copyright.](#)

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.jp/>

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

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



neutone
morpho