Concatenator User Manual

by DataMind Audio

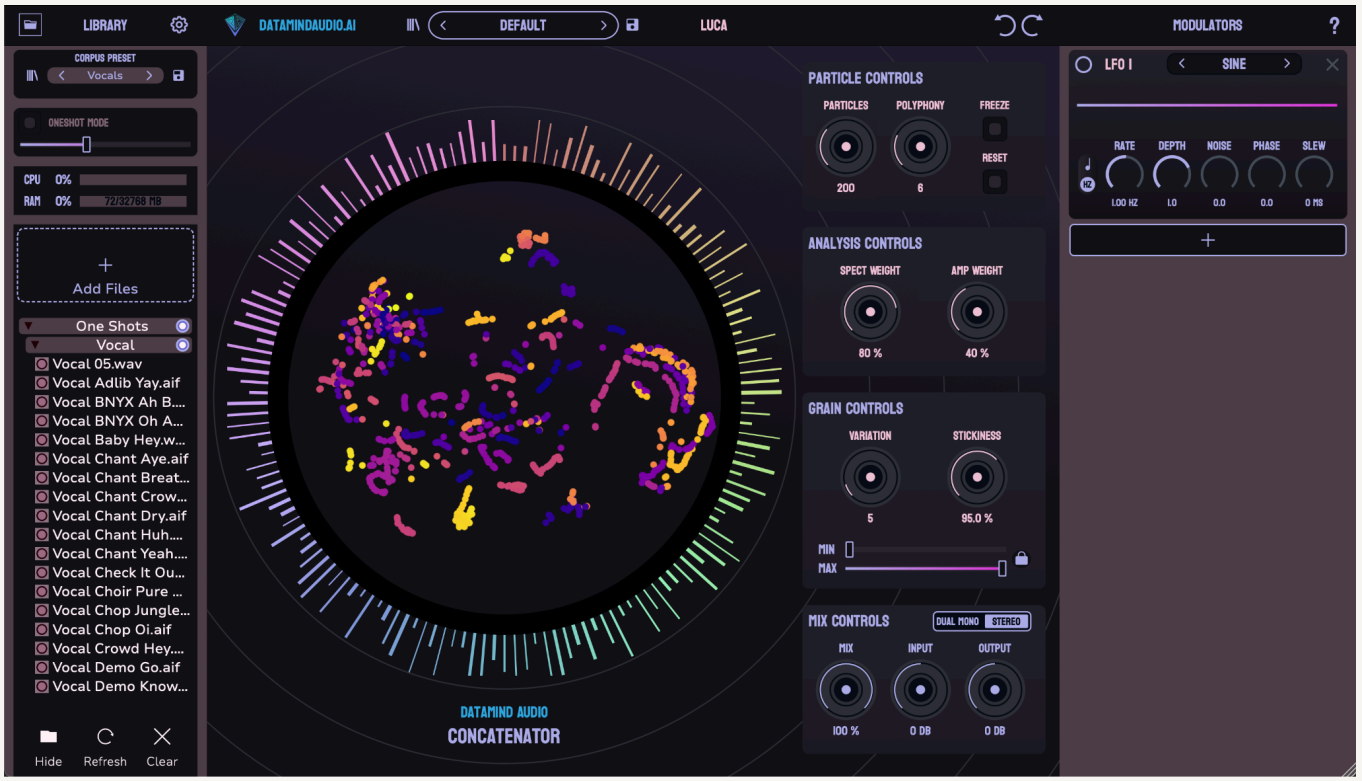

Contents

System Requirements

MINIMUM:
Requires a 64-bit processor and operating system
OS: Windows 10 or later / macOS 10.15 or later
Processor: i5-2500 / AMD FX-6300
Memory: 4 GB RAM
Graphics: 720p monitor resolution
Storage: HDD with 600 MB available space
Additional Notes: Internet connection required for initial activation.

RECOMMENDED:
Requires a 64-bit processor and operating system
OS: Windows 10 or later / macOS 10.15 or later
Processor: Ryzen 5600x / i5-12600 / Apple M1
Memory: 16 GB RAM
Graphics: 1080p monitor resolution
Storage: SSD with 600 MB available space
Additional Notes: Internet connection required for initial activation.

The Concatenator can be installed either as a VST3 or an Audio Unit plug in. All major DAWs are supported (apart from Avid Pro Tools).

Introduction

Welcome and congratulations on purchasing one of the most innovative machine learning based audio tools available.

Concatenator is an AI-powered audio mosaicing plug-in that instantly turns any sound library into an easily playable instrument using any microphone, instrument, or audio file as a real time input. It works by analyzing a live sound input in real-time, and then re-constructing the audio signal using fragments of other sounds.

Concatenator requires two essential components to function:

1. An input audio source, which can be derived from a live microphone, a pre-recorded sound, or a synthesizer.
2. A library of audio files, referred to henceforth as a "corpus." A corpus can consist of a single sound file or potentially hundreds of sounds.

This is digital audio alchemy at its best, enjoy.

Getting Started

To begin using the Concatenator, download and install the latest version from the official website at datamindaudio.ai by following the provided links.

Once the plug-in is installed on your system, launch your preferred Digital Audio Workstation (DAW) and insert the Concatenator into an effects (FX) chain.

DAW Settings

The Concatenator is designed to function across a range of sample rates and block sizes. However, for optimal performance, we recommend configuring your DAW's buffer size to 2048 samples. Smaller buffer sizes may result in decreased plug-in performance, while increasing the buffer size beyond 2048 samples offers no additional benefits.

The ideal sample rates for the Concatenator are 44.1 kHz or 48 kHz. While the plug-in operates effectively at higher sample rates, doing so may reduce latency. Be aware that higher sample rates could also impact the spectral accuracy of the plug-in, potentially leading to less precise pitch tracking.

We are actively working to address these limitations in a future version of the Concatenator.

Toolbar

- **1) Plugin Data Folder**: Opens the plugin data folder in the operating system file browser. Presets and other materials are stored here.
- **2) Settings Menu**: Opens the settings menu. This is described later in the Settings section.
- **3) Info Panel**: Opens a panel that displays information about the plugin build as well as buttons to easily access the log file, website, and user manual.
- **4) Presets Library**: Opens a panel where you can view and select all plugin presets.
- **5) Save Preset**: Opens a panel that allows you to save a plugin preset of your own.
- **6) Username/Register**: Opens a panel that allows you to register the plugin if you haven't already.
- **7) Undo**: Undo the most recent change to the plugin state.
- **8) Redo**: Redo the most recent change to the plugin state.
- **9) User Manual**: Opens the Concatenator user manual in the default operating system PDF viewer.

File Browser

Adding Audio Files

To add audio files to the Concatenator, click the "Add Files" button. A menu will appear, enabling you to browse for a specific folder or individual sound files. Alternatively, you can drag and drop files directly from your file explorer (e.g., Finder or File Explorer) onto the plug-in interface. Once added, these files will load and appear in the file list.

When adding a large number of sound files, please allow time for the analysis algorithm to load and process them. Similarly, if you drag and drop files onto the Concatenator point cloud, a brief processing period is required before they are fully integrated.

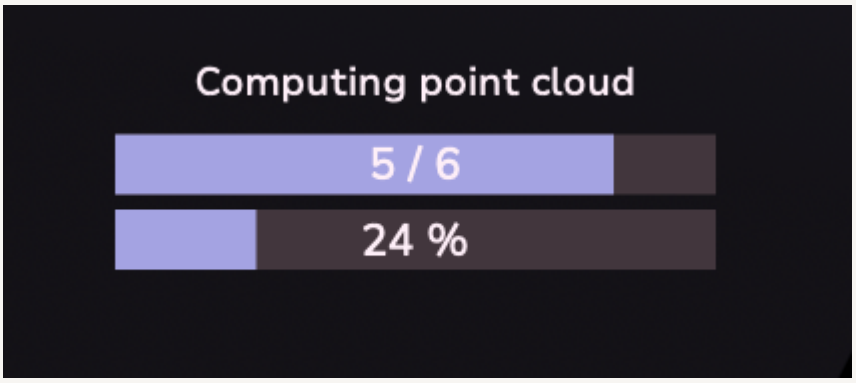

Muting Sounds in Your Corpus

You can mute individual sounds within your corpus by clicking the small button located to the left of each file name. If your corpus contains a substantial number of audio files, muting or unmuting a sound triggers a quick rescan of the library. This update process may take a few seconds as the corpus refreshes.

Navigating the File Browser

The file browser allows you to audition files loaded into the corpus. To preview a file, simply click its name, and playback will begin from the start. To stop playback, click anywhere outside the file browser.

- **Expand or Collapse Folders**: Use the "Hide/Show" button in the bottom left corner to collapse or expand the folder hierarchy.
- **Refresh Contents**: Click the "Refresh" button at the bottom to update the file browser's contents.
- **Clear All Files**: To remove all files from the file browser, click the "Clear" button in the bottom right corner.

CPU and RAM Meters

The CPU meter shows the CPU usage of the current plugin instance as a percentage of the time used in the current audio block. If the audio block size is 50ms and the current instance took 5ms to do its processing, the meter will display 10%. Note that using smaller block sizes in your DAW will cause the CPU usage to appear higher because the audio block size will be reduced but the plugin processing time will remain constant.

The RAM meter shows the RAM usage of the corpus loaded in the current instance of the plugin.

The Point Cloud

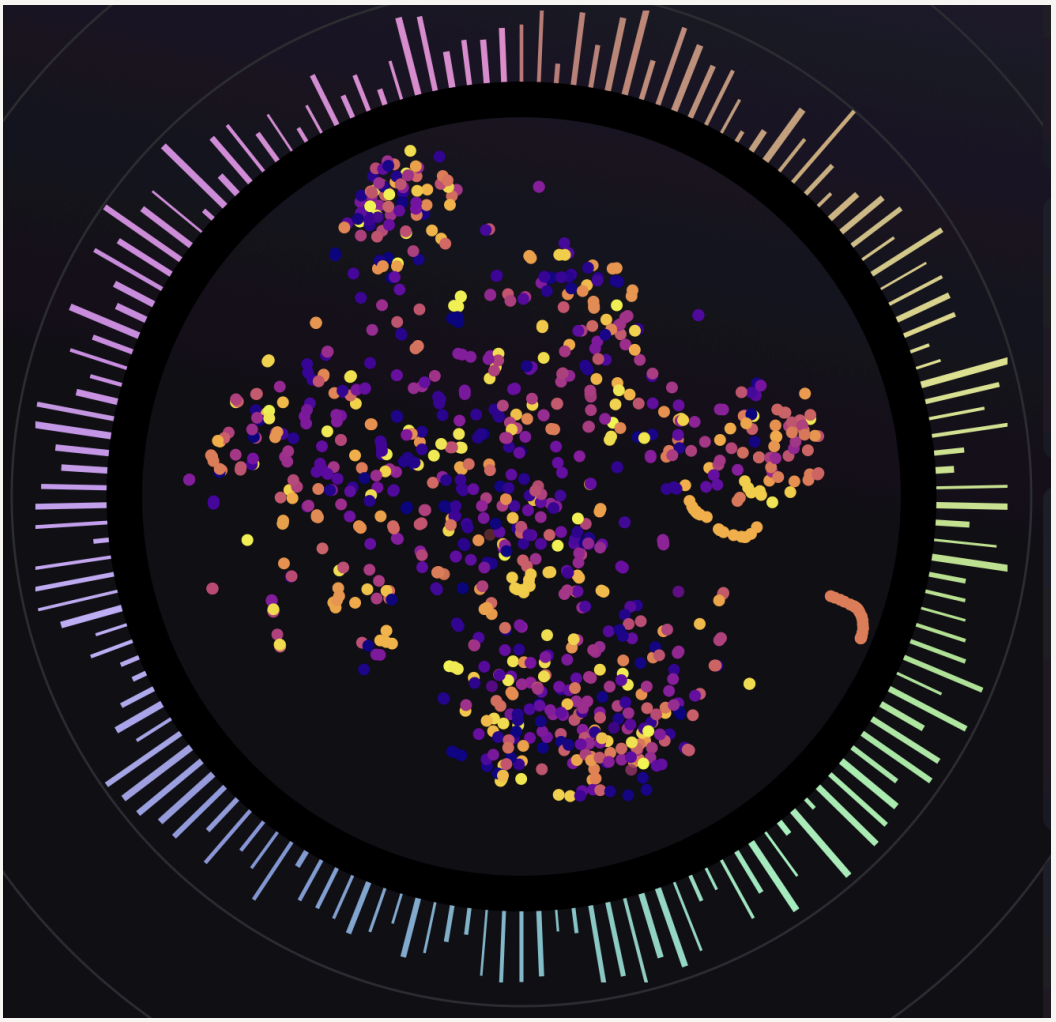

Once the corpus is loaded, the central disk in the Concatenator's plug-in window will display a collection of colorful dots, each representing a cluster of grouped audio grains. As grains are triggered during playback, the point cloud will highlight the clusters containing the currently active grains, providing a visual representation of the audio activity.

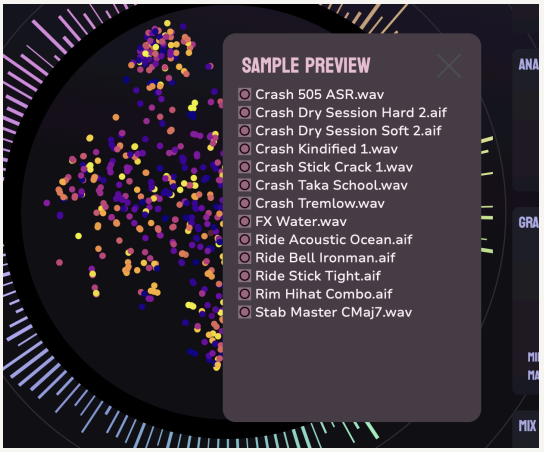

You can interact with the point cloud in two distinct ways:

Trigger Grains Manually: Left-click and drag your mouse over the point cloud to activate various grains, enabling you to audition different segments of your loaded corpus in real time.

Access Cluster Details: Right-click on a cluster to open a small window listing all audio files associated with the grains in that cluster. To preview a file, click its name to initiate playback. You can also disable specific files by clicking the toggle button located to the left of each file name.

Parameters and Controls

Particle Controls

Part of the reason the Concatenator is able to be so much faster than other concatenative synthesisers is that it doesn't try to consider the entire corpus at once. Instead it learns which areas of the corpus are the best match in real time using a machine learning technique called a particle filter.

You can think of the particle filter as throwing lots of darts onto the corpus. Each dart lands on a small slice of audio and those are the slices that will be compared against the input signal to find the best match. The Concatenator uses machine learning to decide in real time which areas of the corpus it should be targeting with each dart.

In the UI we call these darts Particles and there are various controls for changing the behaviour of the particle filter.

Particles

The Particles setting controls the number of particles (or darts) that are active in the particle filter. This governs the computational effort devoted to matching grains to the pitch and rhythm of the input audio. Higher values improve the accuracy of pitch and rhythm reproduction, resulting in a closer match to the input signal, though this comes at the expense of increased CPU load and potential overload. Lower values prioritize efficiency, running faster but with less precise grain matching. Notably, lower settings can yield intriguing, less predictable sonic results—higher isn’t always superior.

Polyphony

Polyphony determines the maximum number of simultaneous "voices" (overlapping grains) that The Concatenator can play from the corpus at any given time. Higher polyphony settings allow for more complex layering of grains, which can enhance the recreation of an input signal but increases CPU usage. Experimentation is key—excessively high polyphony may result in a muddled or "smudgy" sound, while lower settings can produce unique, sparse textures worth exploring.

Freeze

The Particle Freeze button prevents the particle filter from entering what’s called the resampling stage. During particle resampling, particles that are not doing well are discarded and replaced with particles that are doing well. This is a core part of the algorithm that allows the plugin to focus on the best parts of the corpus. However, sometimes you may want the plugin to keep considering some of the poorly performing particles even if it thinks it would be better to focus on something else.

You can think of the Freeze button as a way to stop the plugin from constantly changing the set of audio slices that it is considering and instead just focus on getting the best match out of the audio slices that the particles happen to be currently on.

Reset

The Particle Reset function randomizes the states of all active particles, prompting the plugin to select new grains from the corpus. This can be triggered manually by clicking the Particle Reset button or by pressing 'R' on your keyboard while the plugin interface is active. It's especially effective when combined with high Stickiness and low Spect Weight settings, creating noticeable shifts in the sound. Note that if Spect Weight is set to 0% during a reset, you may encounter "grain splattering"—rapid, chaotic switching between random grains. To resolve this, increase Spect Weight to a non-zero value.

Analysis Controls

This section includes the controls that influence the comparison of audio slices. By changing how the audio slices are compared you can prioritise certain sonic qualities when the plugin is making its decision about what the best particles are.

Spect Weight

The Spectral Weight parameter prioritizes how accurately the pitch and rhythm of the input signal are reflected in the output at any given moment. Increasing this value enhances fidelity to the input's spectral characteristics, though it may cause grains to shift more frequently. Setting it to 0% completely disregards spectral matching, which, when paired with a high Stickiness value (close to 100%), can "freeze" grains for a sustained effect. Be cautious—setting Spect Weight to 0% may lead to unexpected interactions with other parameters, particularly Particle Reset.

Amp Weight

Amp Weight regulates the amplitude smoothness of the output signal. A value approaching 100% ensures more consistent amplitude levels, reducing extreme variations even if the output deviates slightly from the target signal. Lower values allow for greater dynamic fluctuations, potentially aligning more closely with the input's natural amplitude but introducing more pronounced changes. This offers a balance between control and expressiveness.

Grain Controls

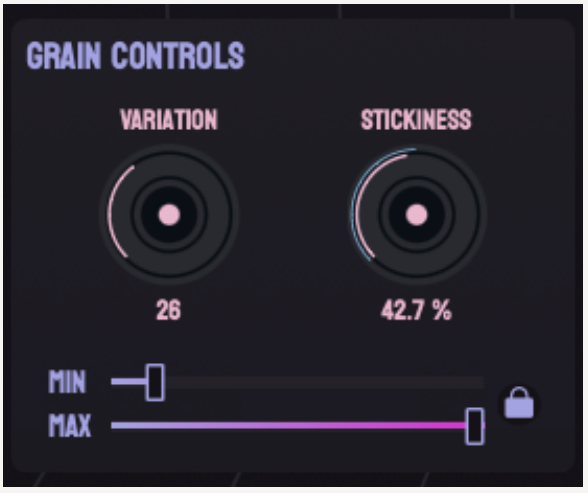

The plugin will often choose multiple time-consecutive slices from the same file in a row. We call these consecutive slices grains and this section of controls allows you to manipulate the behaviour of grain selection and playback.

Variation

The Variation knob adjusts the time interval The Concatenator waits before reusing a specific moment from the corpus for any voice. A higher value increases the delay before repetition, encouraging greater diversity in grain selection. Conversely, a lower value permits

rapid reuse of the same moments, creating a distinctive "stuttering" or glitchy effect. This parameter offers creative control over the temporal behavior of the output.

Stickiness

Stickiness controls the average duration of grains within each voice. A value closer to 100% produces longer, more sustained grains, lending stability to the output. A value near 0% generates shorter, more erratic grains that frequently shift, adding a lively, scattered character to the sound. This parameter significantly shapes the texture and continuity of the grain playback.

Min/Max

In addition to stickiness there are also sliders that allow you to concretely set the minimum and maximum grain lengths. While stickiness is a probabilistic control, the grain length sliders allow you to override stickiness in order to always get a certain minimum or maximum length. By changing the minimum length, you ensure that when a particular grain is chosen to be played, it always plays for at least the minimum length before being replaced by something else. Similarly, the maximum length ensures that when a grain is chosen, once it has played for the maximum length it should be replaced by something else even if the analysis thinks that the grain should keep playing.

You can lock the two sliders together by clicking the lock icon. This will mean moving either slider will also move the other so that the minimum length is always equal to the maximum length. This allows you to ensure that grains are always played for exactly the same amount of time which can give you interesting rhythms.

Do note that currently the grain lengths are limited to a multiple of 1024 samples which is approximately 23ms at a 44.1kHz sample rate. We plan to remove this limitation in a future version of the plugin.

Also note that when the two sliders are locked together any modulation you apply to the minimum length will be discarded and the minimum length will inherit the modulation applied to the maximum length.

Oneshot Mode

When using drums or other percussive sounds in the corpus you may wish to enable oneshot mode to get better matching of percussive events and more reliable transients.

Oneshot mode essentially tells the plugin algorithm that certain files must always be played from the beginning. This may be desirable if the samples you have loaded in the corpus are already carefully split up into short individual sounds e.g. drum oneshots or short sound design clips.

You can enable or disable oneshot mode simply by clicking the toggle button.

Threshold

The slider below the oneshot mode button adjusts the length in seconds that a sample has to be under in order to be considered a oneshot sample. Any samples that are longer than the threshold setting will behave the same regardless of whether oneshot mode is enabled. This allows you to have oneshot mode apply to short samples but still have regular behaviour for longer samples.

Moving the slider all the way to the right gives an infinite threshold, meaning that every sample will be put into oneshot mode regardless of how long it is.

Mix Controls

Stereo Link

Stereo Link toggles between stereo and dual mono processing modes. In stereo mode, the plugin processes the left and right channels together efficiently. Dual mono mode treats each channel independently, doubling the CPU load but allowing for separate processing of the two channels. Choose based on your desired spatial effect and available processing power.

Mix

The Mix knob blends the processed output of The Concatenator with the unprocessed input signal. A value of 100% delivers only the Concatenator’s output, while lower values introduce more of the dry input signal, allowing for subtle or dramatic wet/dry balance adjustments.

Input

The Input parameter adjusts the gain of the incoming signal before it is processed by The Concatenator. This allows you to boost or attenuate the input level to suit the plugin’s response and optimize grain generation.

Output

The Output parameter controls the gain of the processed signal exiting The Concatenator. Adjust this to set the final volume level of the output, ensuring it fits well within your mix or playback system.

Presets

The Concatenator features a robust preset system that allows you to save and load the plugin’s state effortlessly. A collection of factory presets is included to help you get started, which you can explore by clicking the left and right arrows on the preset manager button.

To view all available presets, click the library icon located to the left of the arrow buttons. This opens a window displaying the full preset library. Select a preset by clicking its name to load it instantly. To reset the plug-in to its initial settings, choose the "Default" preset from the list.

To save a new preset, click the floppy disk icon to the right of the preset manager. A window will appear, prompting you to name the preset. Enter a name and click "Save" or press the Enter key to store it. To update or overwrite an existing preset, save a new preset using the same name. Note that the ability to delete presets will be added in a future version of the plug-in.

Important: Presets do not include information about the currently loaded corpus. To save and load corpus configurations, refer to the "Corpus Presets" section below.

Corpus Presets

As you work with the Concatenator, you'll likely want to save specific collections of audio files for easy access later. The "Corpus Presets" feature allows you to do just that. Unlike regular plug-in presets, corpus presets exclusively store information about the currently loaded corpus, enabling you to pair different plug-in presets with various corpus presets seamlessly.

The plug-in includes a selection of stock corpus presets for you to explore upon installation.

To save your own corpus preset, click the floppy disk icon located to the right of the corpus preset manager. This process mirrors the method for saving a regular plug-in preset. After initiating the save, a prompt will ask if you'd like to save an accompanying plug-in preset. If you choose "Yes," a plug-in preset with the same name as the corpus preset will be created, capturing the current plug-in state. Subsequently, loading that corpus preset will automatically load its corresponding plug-in preset, streamlining your workflow.

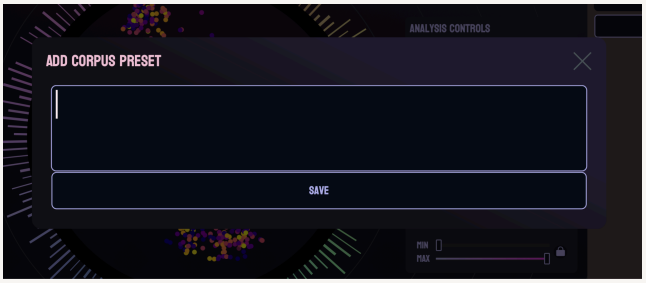

Note: Corpus presets remember which files are loaded based on their file path. If you delete or relocate samples that you have saved in a preset, then that preset will no longer be able to load those samples.

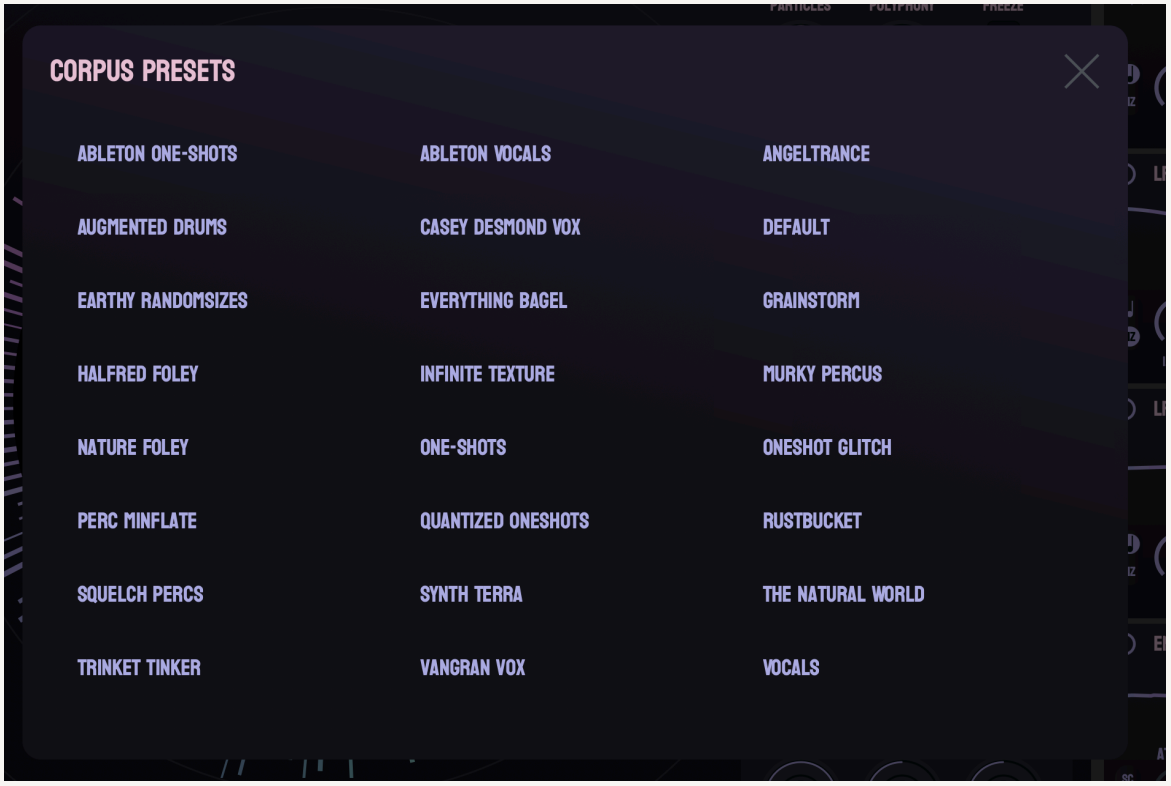

Settings

To access the plug-in settings menu, click the gear icon located in the top-left corner of the Concatenator interface. This menu includes options for adjusting the size of the user interface, which you may need to modify if the interface appears too large for your screen.

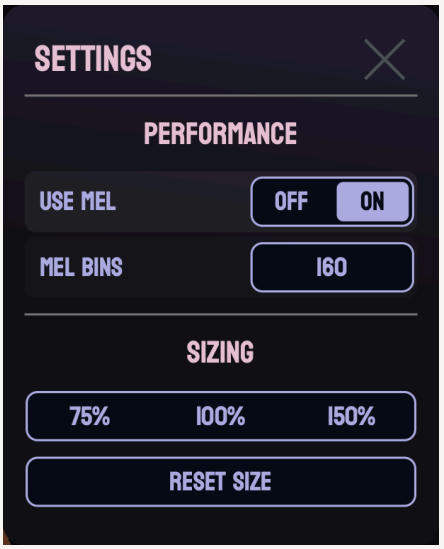

The "Performance" section within the settings menu offers the following options:

Use Mel

The "Use Mel" toggle enables or disables the use of Mel spectrograms. Activating this setting significantly reduces CPU load, making it the recommended default. In rare cases, disabling Mel spectrograms may enhance the accuracy of the plug-in’s output, though this comes at the cost of increased processing demands.

Mel Bins

The "Mel Bins" setting adjusts the number of Mel spectrogram bins used during analysis. Higher values improve pitch tracking precision but increase CPU usage. Adjust this parameter based on your system’s capabilities and the desired balance between performance and accuracy.

Modulators

You can utilize up to five LFOs, five envelope followers, and MIDI envelopes to dynamically modulate parameters within the plugin. To get started, locate the “Modulators” tab and click the "+" sign beneath it. This will display a menu allowing you to select from a total of five LFOs and five envelope followers. After choosing an LFO or envelope follower, it will appear on the right-hand side of the interface. Assigning a modulator to a parameter is straightforward: simply drag the handle from the small dot in the top-left corner of the modulator to the center of the parameter you wish to control.

LFO:

- **LFO Rate**: Adjusts the speed of the low-frequency oscillator’s cycle.
- **LFO Depth**: Controls the intensity or amplitude of the LFO’s effect.
- **LFO Noise**: Adds random variation to the LFO waveform for a less predictable modulation.
- **LFO Phase**: Sets the starting point of the LFO cycle for timing adjustments.
- **LFO Slew**: Smooths transitions in the LFO signal, reducing abrupt changes.

Envelope Follower:

- **Envelope Follower Attack**: Determines how quickly the envelope responds to an input signal’s onset.
- **Envelope Follower Release**: Sets the time it takes for the envelope to fade after the input signal drops.
- **Envelope Follower Depth**: Adjusts the strength of the envelope’s modulation effect.
- **Envelope Follower Sidechain**: The SC/IN button toggles the audio input to the envelope follower between the main audio input to the plugin and a sidechain input that you can feed in from your DAW. For more information about setting up audio sidechains, please refer to the user manual of your DAW.

MIDI Envelope:

- **MIDI Envelope Attack**: Defines how fast the envelope rises after a MIDI note is triggered.
- **MIDI Envelope Decay**: Sets the time it takes for the envelope to fall from its peak to the sustain level.
- **MIDI Envelope Sustain**: Adjusts the level maintained while a MIDI note is held.
- **MIDI Envelope Release**: Controls how long the envelope lingers after a MIDI note is released.
- **MIDI Envelope Velocity**: Scales the envelope’s intensity based on how hard a note is played.
- **MIDI Envelope Depth**: Determines the overall impact of the MIDI envelope on the modulated parameter.

Modulation

To utilize the modulators in the "Modulators" section, locate the rings in the top-left corner of a modulator. Hover over a ring, then click and drag it toward a parameter. As you drag the modulation handle, all modulatable parameters will be highlighted, as shown in the accompanying image. Release the handle over a highlighted parameter to link it to the modulator.

Once connected, small colorful dials will appear atop the modulated parameter. Adjust these dials to control the amount of modulation sent from the source to the destination parameter. Each dial operates within a range of [-1, 1]:

- **1**: Sends 100% of the source signal to the destination parameter.
- **-1**: Sends 100% of the inverted source signal to the destination parameter.
- **0**: Sends none of the source signal to the destination parameter.

To disconnect a modulation source, double-click the corresponding dial.

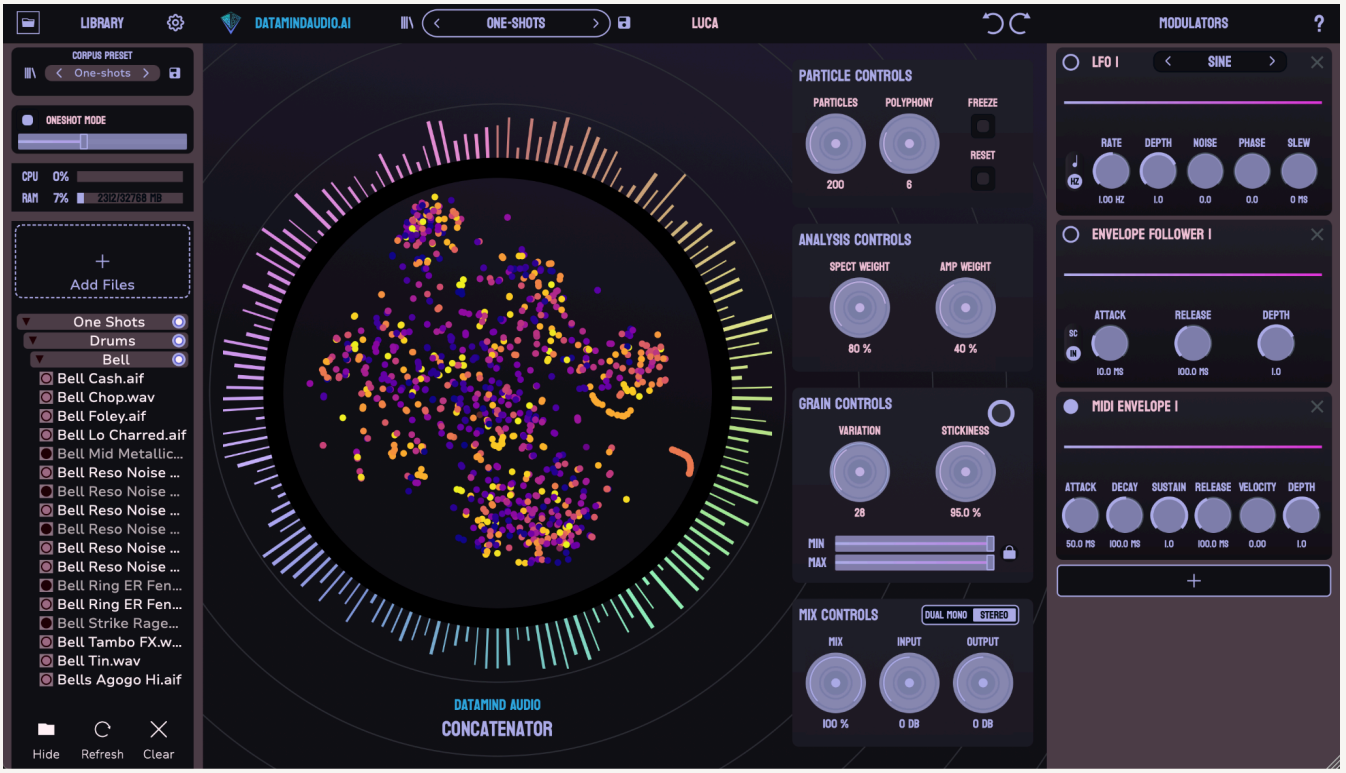

CAUTION if applying modulation to polyphony and to particles parameters as these are directly linked to CPU load.

Getting Help

DataMind Audio is here to help you. Join our discord to reach our team with queries. You can also email us, but by joining our Discord channel you have direct contact with the developers working on the plugin!

Discord: discord.gg/x4AKrZZVmd
Info email: info@datamindaudio.com
Support email: support@datamindaudio.com

P.S. Information about the plug-in build, directories used by the app and log files can be found by clicking on the datamindaudio.ai button in the top-left. If you encounter any issues and would like to file a bug report, please email us or post a message on our Discord bug-reports channel and attach the Log.txt file that can be found in the Logs folder. Thank you!

Note: The Log.txt file may contain the paths of files that have been added to the corpus. If you are working on something super secret and don't want to share the log file on a public channel, don't hesitate to email us or DM either Robin or Luca on Discord.

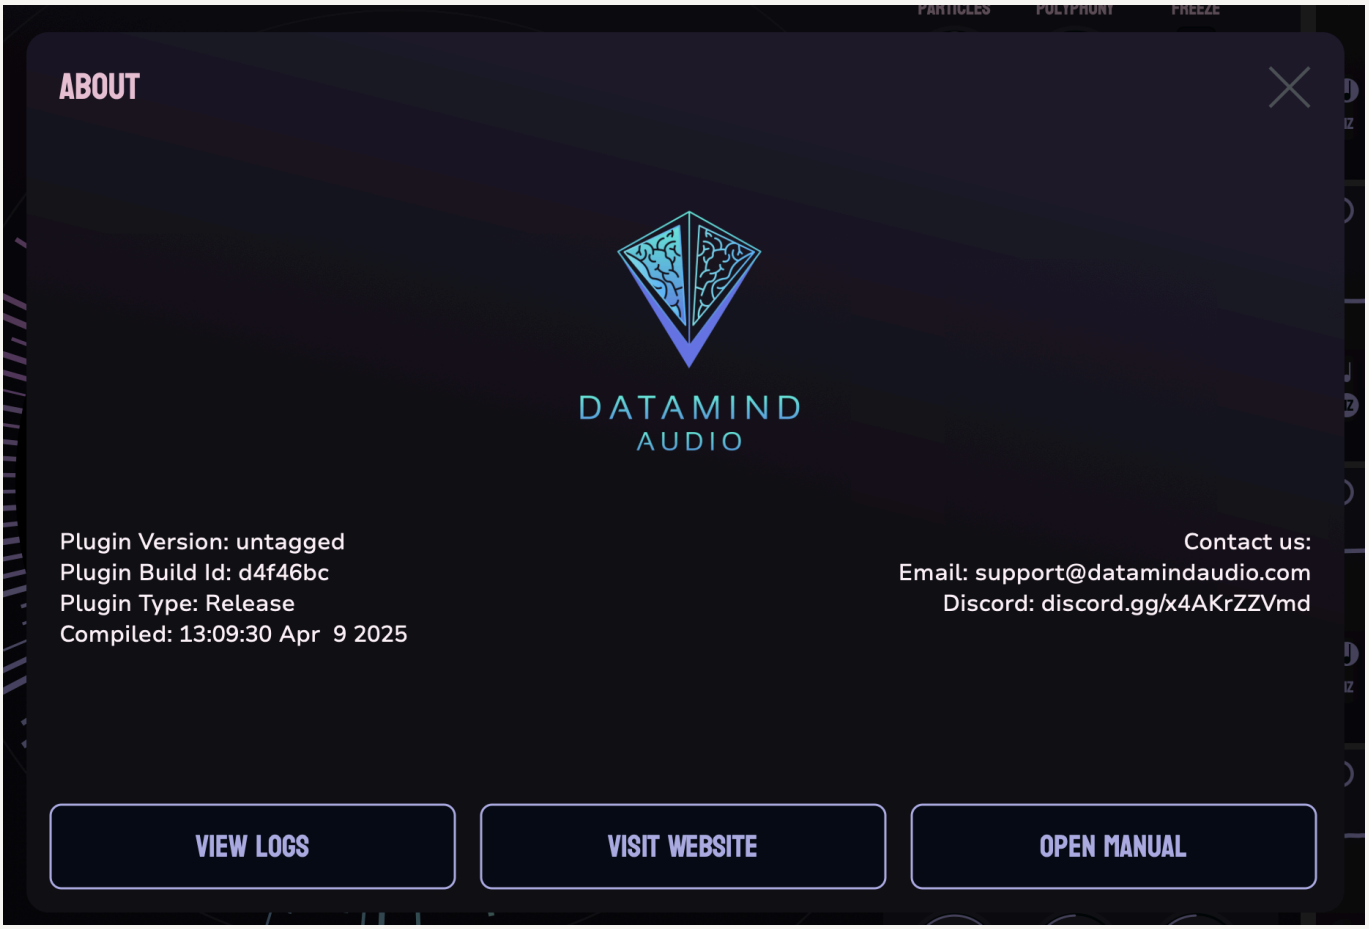

Credits

Dr. Christopher Tralie - Lead Research Consultant and Algorithm Inventor
Ben Cantil - Product Designer
Robin Leathart - Lead Programmer
Jacob Crider - UI/UX Designer
Luca Ayscough - Programmer
Yashique Chalil - Sound Designer

“Vox Casey Desmond” Recorded and Edited by Casey Desmond (Soundcloud / IG), All Rights Reserved

“Vox vangran” Recorded and Edited by Vanessa Graniero (aka vangran) (https://vanessagraniero.com/) and George Sideris, All Rights Reserved

“Halfred Foley” Recorded and Edited by Alfredo Treccani (aka Halfred / Halfreak) (Soundcloud / Bandcamp / IG), All Rights Reserved

Special thanks:

Rob Clouth (https://robclouth.com)
Martin Parker (https://tinpark.com)
Nolan Petruska (https://frequent.audio)

“self-ish” presets by _self (Soundcloud)
“Viberous” presets by Viberous (Soundcloud)