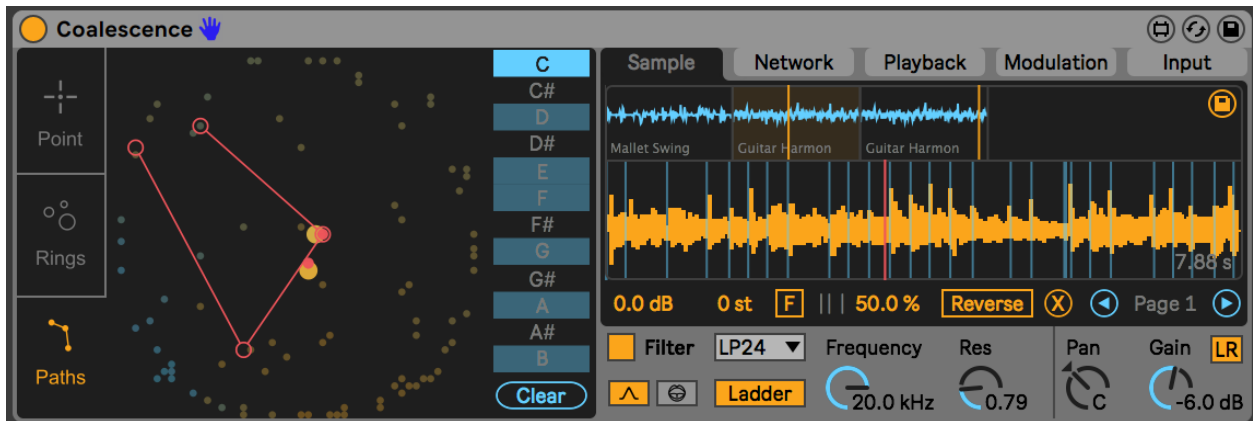


Coalescence User Manual




To install:

Unzip the folder and drop the folder called “Coalescence” in this EXACT location in order for presets and samples to load correctly: `../ableton/user library/presets/instruments/max instrument`
For best results or if you are having issues, make sure you are using the latest version of max/msp. You do not need to have a license if you are using Live suite. Download the newest version here: <https://cycling74.com/downloads> and once downloaded go to the ableton Preferences > Library and set the newly downloaded version of max to the one ableton should use. (Also you can try to see if it works fine with your bundled version first).

Synopsis

Coalescence is a MaxForLive Instrument that combines concatenative sampling together with a neural network. Coalescence takes multiple samples (up to 2000) and analyzes them and feeds slices of the samples into the neural network (an SOM) which organizes the samples into a 2D map of cells, which you can see visualized as little circles. The device plays around with this organization of sound slices using three different playback modes: Point, Rings, and Paths.

The Network

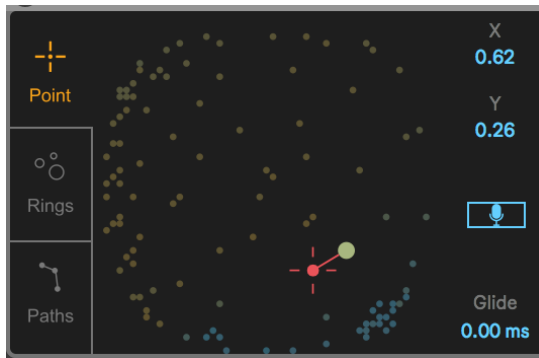


The SOM network has 2500 cells which can “hold” sound slices. This is good to remember because even though you can drop up to 2000 samples, the cells of the network may fill longer before you reach that number. The network organizes slices into clusters based on the chosen spectral features to analyze (spectral features being a description of the slice’s frequency content). It will cluster the slices based on their similarity to each other and therefore organize them meaningfully in a 2D space (such as in similar tonality, timbre, etc)..

The slices are visualized as colored dots (Note: hold alt/option while hovering over the network to highlight all of the slices). The colors are determined based on the low/high intensity values of

each slice's spectral features. The type of feature the network is trained on is determined in the Network settings tab, which I'll get into later. The device chooses which sample to playback and the starting point through traversing this map.


Point Play Mode




This playback mode works the closest to a classic sampler. In this mode you choose a single starting point/sample to playback from. MIDI notes trigger a voice and repitch it based on the distance from middle C (C3) just like a normal sampler. The lookup position can be chosen by dragging this mouse on the network or by changing the X and Y parameters. There is a *Glide* parameter for changes in the position. If the *Mic/Input* button is enabled, the position will be controlled by a routed

incoming audio which is handled in the Input settings (described later). How that works is that for every moment in time (or spectral frame of 1024 samples) the routed incoming audio is sent to the network and the network gives the 2D coordinate on the map that most closely matches the incoming audio. Therefore, if a sample slice in the network is similar to the incoming audio then the 2D position will go near or at that slice. This can be a powerful tool for many experimental uses!

In the Playback settings tab on the far right are settings for the different play modes. For Point mode there are a few parameters. *Retrigger* allows for one-shot (non-looping) voices to be retriggered after their playback has ended if the current sample slice changes. *Note Glide* allows for portamento pitch gliding between voices. *Latch* allows for the chosen slice to not be changed after the voice starts playing (Note: modulations will still move it!)




Rings Play Mode



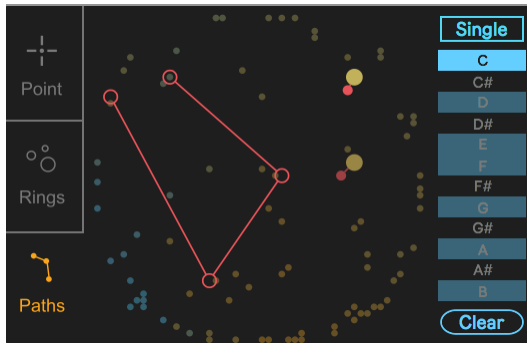
In Rings mode incoming MIDI notes trigger slices within circular ranges called rings. Each ring has an assigned MIDI note that triggers it. When triggered, a randomly chosen slice within the ring's radius will be played. Rings can be created and managed *Manually* or *Automatically*. If manual, you can create a ring by double clicking the network, drag to move them, control/command + drag to resize the rings, and shift + click to delete a ring. Each

ring created manually will automatically assign in a MIDI note pitch starting from the *Start* pitch and going up. You can also change the assigned note by clicking a ring and changing the *Current* parameter. In *auto* mode rings are created and spaced automatically based on the *# Auto* parameter and cluster detection algorithm. Resizing a single ring will change the radius for all the rings in *Auto* mode. The notes are automatically assigned from left to right to reflect a keyboard starting from the *Start* pitch. Changing the *Start* pitch will shift all of the assigned notes. Pressing *Reorder* in Manual mode will reassign all of the notes from left to right like in *Auto* mode. *Clear* clears all the rings.



These are the playback parameters for Rings. *New Slice Retrigger* works exactly as *Retrigger* in Point mode. *Selection* parameter has two values: *Range* and *Center*. If in center, then when a Ring is triggered it will only play the slice in the center of the ring instead of picking a random slice within its range.


Paths Play Mode



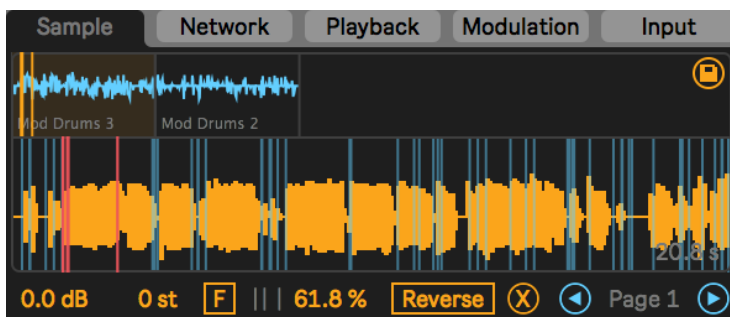
In this mode you can create playback paths that can be assigned to different pitches (octaves are irrelevant). In this example a path is assigned to “C” and whenever a MIDI note of C (of any octave) is received this path will be triggered for playback. To create points in a path double click on the network, drag to move points, shift + click to delete them. *Clear* clears the paths of a selected note.

Additionally you can put Paths into *Single* mode, which makes it so there is only one path that is triggered by any MIDI note and in this mode MIDI pitches instead repitch the playback like a classic sampler.

The playback parameters for path are as follows: On top is the *Loop* setting on the path playback. Left is one-shot, middle is loop, right is boomerang. The *Reverse* option reverses playback. The *Interval* sets the time it takes to travel between points. *Slice* has two options: *Path* and *Points*. When set to points, then the path playback only jumps to the points instead of moving along the lines of each path.




Sample Settings



Here you can manage the samples. You can drop samples or folders of samples on the top multi-sample

selector display (Note: you can also drop them on the network itself). You can select the sample for editing on the top selector, and on the bottom right you can move between *Pages* of samples with the arrows. The *Save* icon at the far right if enabled saves analysis files for the dropped samples so you don't have to wait for them to be analyzed in the future which is create for preset or set loading and also for long samples you use in the device constantly (Warning: this will create .dbd analysis files in the same location as the sample, similar to Ableton's .asd!). Below this you can the selected sample display. You can drop individual samples on this display to switch the sample at that slot. The blue lines are the transient slices (in Transient mode only). You can drag the mouse on the opened sample to select the range of audio you want to send into the neural network. Holding shift + dragging on the sample previews in the network where each slice is. Below the display are the individual sample parameters (they only affect the opened sample). You can change its volume, pitch and fine tuning, the transient sensitivity, reverse the sample, and remove the sample. Holding shift while changing any sample parameter such as volume, pitch, direction, etc will change that parameter for every sample!

Network Settings



Here is where you manage the SOM neural network. On the left are parameters for training the network. There is the number of *Repetitions* that the network does each time a new training is triggered. Each repetition every transient or spectral frame of the sample is fed into the network. *Strength* sets the amount the network

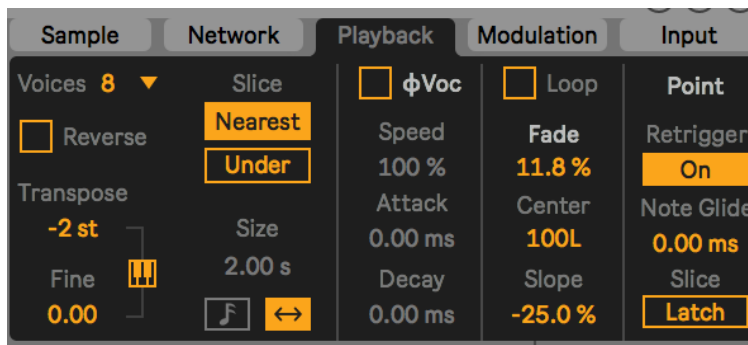
learns each repetition. Higher strengths cause the network to learn and cluster faster, but it will also become biased to data faster. *Restart* clears and resets the network and starts training every sample from scratch. *More* continues further training without clearing the network. *Stop* is an emergency stop (note it can cause inaccuracies, but it is good for a hard stop if the training is taking too long). *Append* if enabled causes for changes in samples or dropping of new samples to only train the affected sample and not the others (I recommend keeping it on!).

The middle section deals with the network settings. *Features* tells the network which spectral feature of the samples to train on. The options are... *Chroma*: this will analyze based on the chromatic tonality of the sound (based on a 24 step octave), it is good for organizing sounds based on their tone such as for melodic or similar purposes. *Mel* and *Bark*: these two modes are filterbanks of the spectral features with scalings informed by psychoacoustic studies of bands with perceived equal steps. These modes are good for organizing slices based on their frequency content from low to high pitch (such as for percussions, etc) and differ from each other in scaling so will yield different results. *Speech*: is actual the mel cepstral coefficients (MFCC) which are commonly used for speech analysis. Use this mode to group slices based on

speech such as vowels. *Radius* changes the size of the sample slice clusters formed by the network. This is something to dial in based on the number of slices and grouping types desired. For example if there are only a few groups you'd like you can have a higher cluster radius. However, smaller cluster sizes will populate the map much slower, so you could start training with higher ones and then shrink. *Transients Only* if enabled only transients will be sent into the network as slices. If this is off the EVERY SPECTRAL FRAME will be sent into the network. This will cause much much longer training times especially for long and many samples! Therefore I usually recommend to have transients on and you can turn up the samples' transient sensitivity to get more slices. However, often times you will want many slices from a sample to glide between them and therefore may want to turn this off. *Clear* clears all of the samples and the neural network (Use with Caution cannot be undone!)

The right section allows for a preview playback to occur whenever you click or drage the mouse on the network. You can control the preview's *Gain* and loop *Length*. You can also have the preview only be the slices directly underneath the mouse or the *Nearest* slice to it.

Playback Settings



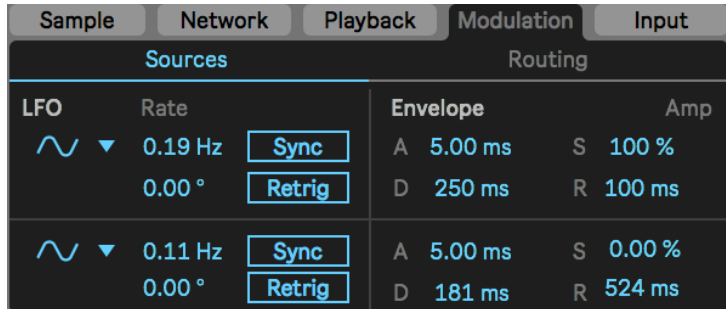
Here you control the sample playback. You can set the number of poly *Voices*, *Reverse* the playback, and *Transpose* and *Fine* tune the pitch. The *MIDI Keyboard* button if enabled allows the MIDI notes to repitch the playback like a classic sampler (for Point mode only). *Slice* has two options:

Nearest and *Under*. Nearest allows

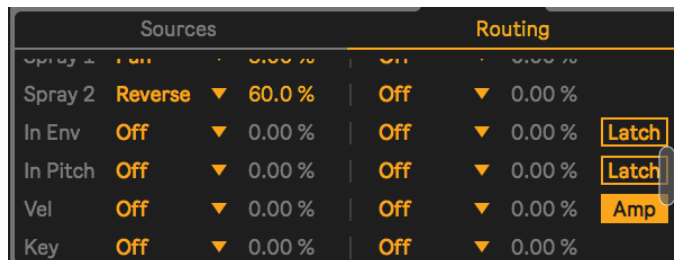
for the nearest sample slice to the chosen cell to be played back. If under is selected, then it only plays the slice directly in the cell chosen and otherwise no playback occurs. *Size* sets the playback length/loop size and it has synced interval options. The <-> option for size allows the sized to be based on the transient length (for Transient mode only) To the right is a Phase Vocoder playback mode . This mode decouples pitch from speed and allows time stretching by playing back spectra in the frequency domain. It does have a slushy/spectral quality to it! The *Attack* and *Decay* parameters can be used to interpolate between spectra changes. This allows for a few (albeit “cheap” sounding) effects. *Attack* allows for a kind of fading between sounds effect which can be useful for fading between changes in slices. *Decay* creates a cheap spectral blur or reverb-like affect, and attack and decay can be dialed in in tandem to support each other. To the right of this is a *Loop* playback option and below it is a *Fade* window which can be used to remove looping clicks but also to shape the grain. If *Center* and *Slope* are at 0% and *Fde* at 100%, the shape of the window is a triangle. Moving the center left to right changes this to a saw down/up and the *Slope* adds an exponential slope if positive and log slope if negative.

NOTE: if doing a very low speed playback in Phase Vocoder then you should remove the Fade window (set fade to 0%) as it will freeze at the beginning of the fade window and be silent.

Modulation Settings



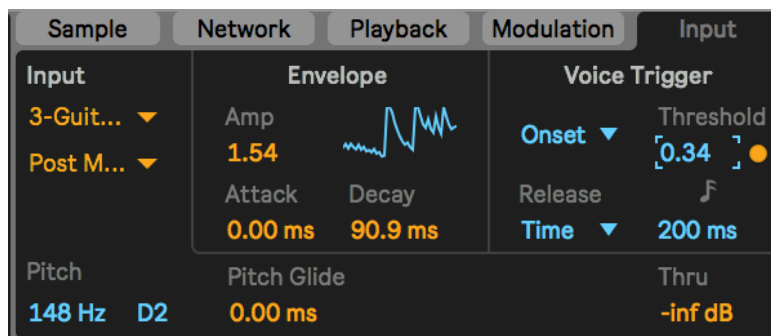
There are two sub tabs in the Modulation tab. The left one is for the *Source* settings. There are two LFOs and two Envelopes with traditional settings (The LFO's bottom shape option is perlin noise). The top envelope is also an Amplitude envelope by default.



The *Routing* tab is where you map modulators to destinations. Each modulator has 2 destination options each with a bipolar attenuator. In addition to the two LFOs and envelopes, there are two random *Spray* values (which are create at the beginning of each voice), the input audio's envelope follower and pitch

detector (described in the Input section below), and standard MIDI values such as Pitch Bend, velocity, aftertouch, etc. The input modulators have *Latch* options which latch the value (sample and hold) at the beginning of a voice. The velocity modulator has an *Amp* button which chooses if the velocity controls the amplitude of the voice.

Input Settings

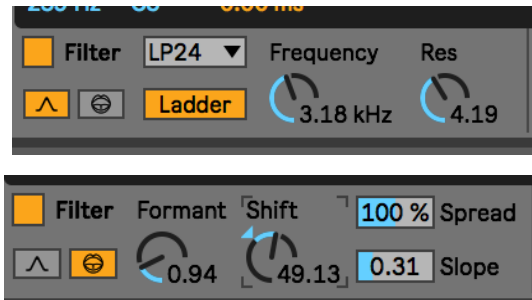


Here you can *Route* an input audio from another track for various uses which have been described earlier. As mentioned in the previous section, the input provides a pitch detector and envelope follower. You can change the *Glide* of the (rudimentary) pitch detector and

the envelope and volume of the follower which can both be used as modulation sources. There is also a *Thru* gain for the audio on the bottom right. The top right section allows you to have the incoming audio transients trigger voices instead of the MIDI keyboard. It will trigger voices at the last received MIDI pitch. It has a transient *threshold* setting and a *Release* setting. In *Time*


release the voices will be released after a set interval. In *Legato* release it holds the voice until a new transient trigger is received. Look into the Point Mode section for further uses of the input.

Filter



There are several filter options. By default is a simple biquadratic filter, with a *Ladder* filter options and shapes. Additionally, there is a *Formant* filter mode in which you can select the *Formant* (a tenor-alto blend I made), *Shift* their center frequencies, adjust their *Spread* and *Slope* adjusts their bandwidth.

Mix



Here you have the master panning and gain. The *LR* button enables a stereo mode processing for stereo samples. This doubles the amount of filters, etc for each voice so it is slightly more efficient with it off (in which case stereo samples are summed to mono).

I hope you enjoy this device! Please email me if you have bugs or other issues:

dillonbastan@gmail.com

More: <http://dillonbastan.com>