



FACTORSYNTH

user manual

www.jjburred.com - software@jjburred.com
© J.J. Burred, 2018-2021
factorysynth version: 2.1, 12/2/2021



Installation	2
Workflow overview	2
Interface	3
Main control panel	5
Element mixing and editing	8
Analysis panel	9
Clear panel	10
Export panel	10
Mixer panel	11
Working with presets	11
Automation and MIDI-mapping	11
Usage from Max	12
Factorization how-to	13
From Factorsynth 1 to Factorsynth 2	13
Support	14

Factorsynth is a Max For Live device that uses a machine learning technique (matrix factorization) to decompose sounds into sets of elements. Once these elements have been obtained, you can modify and rearrange them to modify your audio clips in many different ways. The obtained elements can represent anything from individual notes or drum hits to attack portions, background noises, drones, repetitive patterns, and so on, and they can be overlapping in the original sound, in both time and frequency.

Here are some examples of things you can do with Factorsynth:

- Remove certain elements (notes, drum hits) from clips
- Change the rhythm of a drum loop by separately displacing its components (kick drum, snare, etc)
- Randomize the timbre and the internal temporal structure of sounds
- Create infinite rhythm and timbre variations of your loops while mixing in session view
- Create complex textures out of a song excerpt
- Create new elements consistent with the overall timbre and rhythm of the original sound
- Upmix mono clips to stereo clips by panning the individual components
- Apply audio effects selectively only to certain elements of the sound (e.g. apply reverb only to the snare)
- Separately process attack and sustain portions of notes, or consonants in voice signals.
- Obtain hybrid sounds by using a new kind of component cross-synthesis (e.g. making each drum element drive a separate timbre element of a second sound)
- Discover hidden components in sounds that you didn't notice before

Thanks for purchasing Factorsynth, and happy factorizing!

Installation

To start using Factorsynth, just drag and drop the factorsynth.amxd file into the audio effects area of an audio track in Ableton Live 9 or 10. Once the device has been loaded, you should see the small device window pictured on the right. Click on the logo in that window to open Factorsynth's main interface.

If you're installing an updated version of Factorsynth, it is recommended that you delete the old factorsynth.amxd file.

For a quick start guide, click on the  icon on the upper part of the main interface.



Interface zoom: The zoom control on the device window lets you zoom in/out the main interface window.

Workflow overview

Unlike traditional audio effect devices, which take the track's audio as input and generate output in real time, Factorsynth is a **clip-based** device. It works on audio clips from your Live set that have been loaded into Factorsynth by drag and drop. Once an audio clip has been loaded into Factorsynth, it will be decomposed into elements (exactly how to do this will be covered in next sections). The decomposition process is called **factorization**, because it is based on a technique called non-negative matrix factorization (NMF).

Factorization usually takes a few seconds, and can be performed while the Live set is playing. Once the factorization is ready, you can playback and modify your sound by modifying or recombining the extracted elements. Playback of the processed sound is controlled by Factorsynth's own playback button.

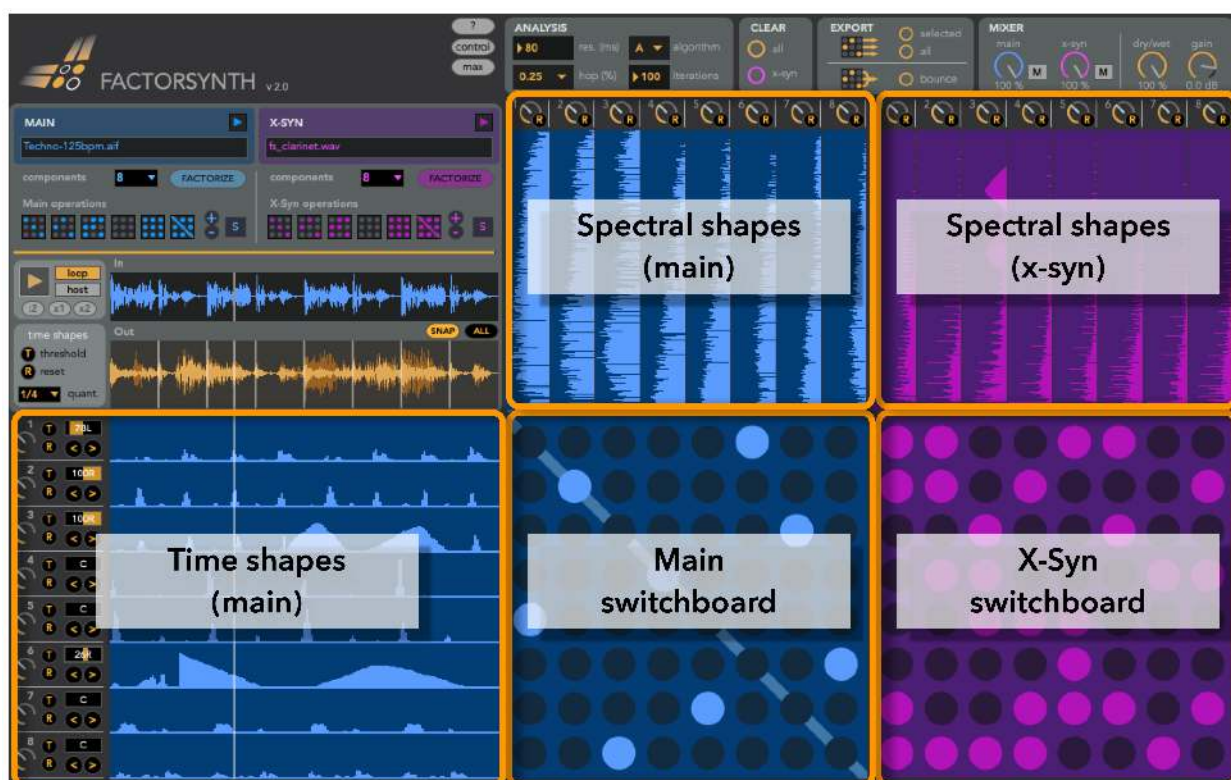
Since it is a clip-based device, Factorsynth will only affect the clip that is currently loaded into the device, even if the track contains other clips. Also, since the device takes its input audio from the loaded clip, it must always be in the leftmost position of an audio effects chain: it will ignore any processing happening before it on the effects chain (you can of course process its output with any other audio device).

The clip that you first load into Factorsynth becomes the **main sound**. It is the sound that will provide both temporal and spectral elements (more on this below). You can optionally load a second sound, called the **x-syn sound** (for "cross-synthesis"). This sound is used to add new spectral elements to the palette for sound creation.

The original main clip's "transpose" and "detune" settings from Live's clips editing window will affect Factorsynth's output in real time. However, changing the original clip's gain and editing local warp points is not supported and won't have any effect on Factorsynth's output.

Interface

Controls on the main interface are color-coded: blue controls correspond to the **main sound** and purple controls correspond to the **x-syn sound**.



The main interface is organized around several control panels and the following sections:

Time shapes

Factorizing a sound creates two types of elements: time (or temporal) shapes and spectral shapes. Temporal shapes are displayed on the lower left of the interface. They capture the temporal evolution of distinct sound events within the original sound. For example, if you are factorizing a drum loop, one of the temporal shapes will likely contain peaks that correspond to hits of the hi-hat, another one will contain peaks when the kick drum plays, and so on. In Factorsynth, temporal shapes are always provided by the main sound.

Spectral shapes

They can be understood as the “timbre ingredients” that make up the original sound. They are fixed spectra, displayed vertically, with lower frequencies on the lower part and higher frequencies on the upper part. Spectral shapes can be provided by the main sound (blue spectra) or by the x-syn sound (purple spectra).

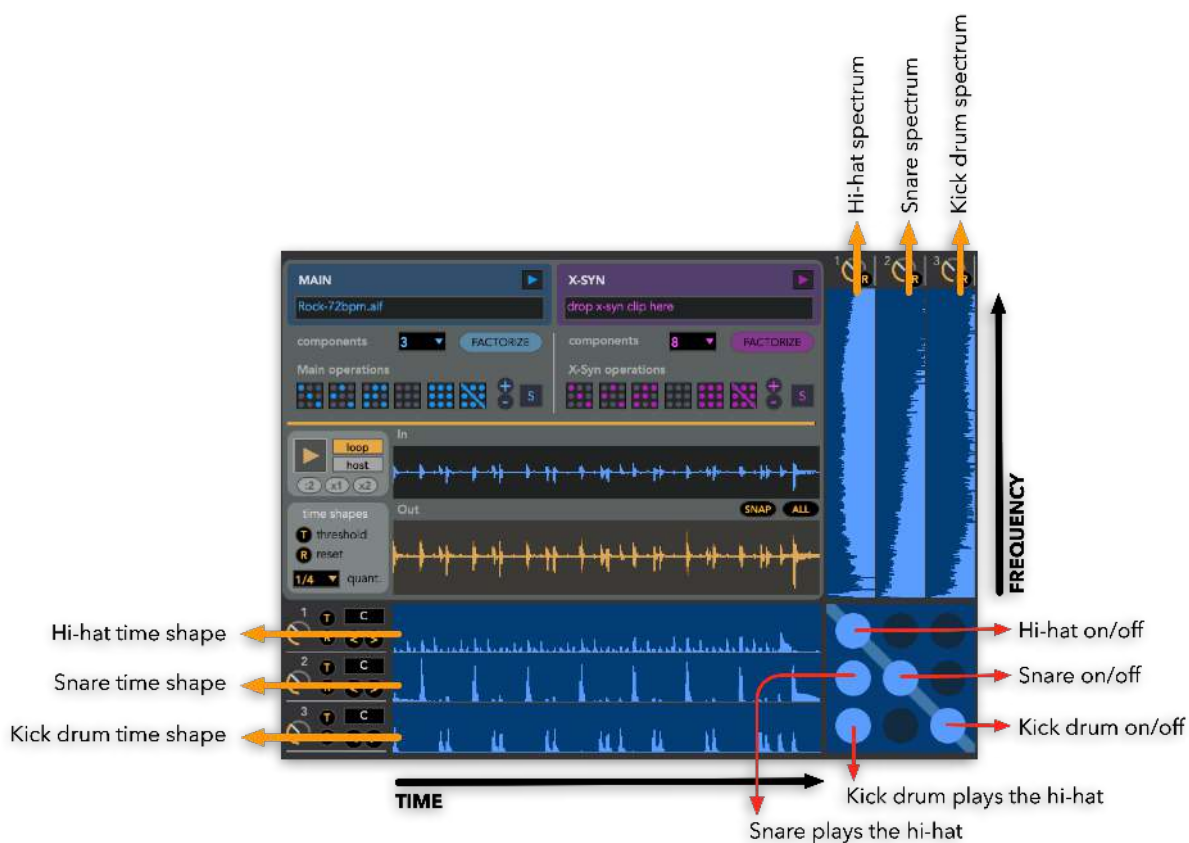
For example, in the image above, the main sound is a drum loop. It's spectra (blue) are thus wide-band and noisy. The x-syn sound is a clarinet sound, and thus it is possible to see many harmonics as horizontal lines on the purple spectra.

Spectral shapes are sorted by “brightness”: the first (leftmost) spectral shapes are the richest in high frequencies. For instance, in a drum loop, the spectral shape corresponding to the hi-hat will likely be the first, and the one corresponding to the kick drum will likely be one of the last.

Switchboards

Sound output is created in Factorsynth when you connect temporal shapes with spectral shapes. These connections are created by the buttons on the two switchboards: the main switchboard connects time and spectral shapes of the main sound, and the x-syn switchboard connects time shapes of the main sound with spectral shapes of the x-syn sound. The connection of a time shape with a spectral shape is called a **component**. The key in understanding how Factorsynth produces sound is the following: a component is created by a time shape that *modulates* a spectral component. In other words: a component is a fixed spectrum that changes amplitude in time according to the temporal shape, which acts as a modulator. By setting a switchboard button on/off you are adding/removing components to/from the output sound.

You'll notice that the **diagonal** is highlighted on the main switchboard. This is because main components on the diagonal are special: they are the original components present in the input sound. In contrast, **off-diagonal**



components are new components that were not initially in the sound, created by combining originally unrelated shapes. Let's look at a simple example to illustrate this.

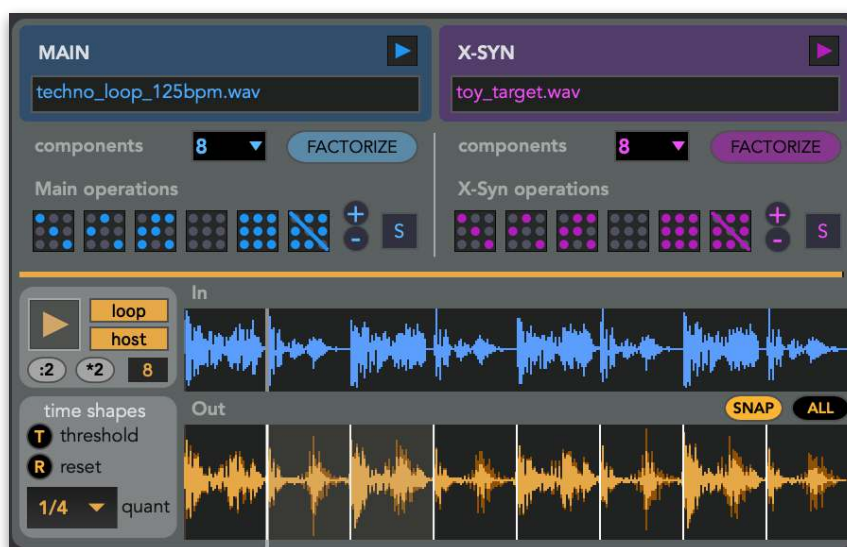
The figure above shows the factorization of a simple drum loop into 3 components: the hi-hat, the snare and the kick drum. Clicking on the components on the diagonal, you will add/remove the corresponding original sounds of the hi-hat, snare and kick drum. Clicking on off-diagonal components, you will generate new sound events that were not present in the original drum loop: the figure shows two examples: selecting the lower-left component makes the temporal pattern of the kick drum play the timbre of the hi-hat. Selecting the component just above it makes the temporal pattern of the snare play the hi-hat sound.

Creating connections on the x-syn switchboard makes temporal patterns of the main sound control spectral shapes of the x-syn sound. This is a special form of *cross-synthesis*, hence the name x-syn. Playing with the x-syn switchboard, you can do things like making each instrument from a drum kit (main sound) drive different piano sounds (x-syn).

Switchboard connections can also be created automatically via the operation buttons on the control panel (see below).

Main control panel


The main control panel manages sound loading and playback, launches factorizations, and creates automatic switchboard connections. The left section (blue) controls the main sound, the right section (purple) controls the x-syn sound.




Factorizing the main sound

The first step in using Factorsynth is to factorize a main sound. In order to do that:

1. Select the desired number of factorization components (2 to 30) from the blue component drop-down menu. For the main sound, this is the number of temporal shapes and spectral shapes that will be obtained. This parameter has a strong impact on the sound results (see section "Factorization how-to"). The default value is 8 components.
2. Drag and drop any audio clip from your Live set into the blue area to the upper left. Factorization will start right away, as indicated by the yellow progress bar. The name of the loaded main clip will be



displayed on the blue text window.


 You can also drag and drop sound files from Live's file browser or from your operating system's file browser. However this can produce looping problems, so it is recommended you always import the clip into the Live set, then from there into Factorsynth.

3. Once factorization is done, the input waveform and the main temporal and spectral shapes will be displayed. The switchboard will still be empty, so you will now need to create components to start hearing any sounds. You can do this either by clicking on any connection buttons on the main switchboard, or by clicking one of the **operation buttons** (see below) that will automatically create connections on the switchboard.

You may want to re-factorize the same clip after changing the number of components or the analysis parameters. In that case, you don't need to drag and drop the clip again, you can just click on the factorize button.



When re-factorizing the same clip, Factorsynth's old output will keep playing until a new component is created or modified. However, when factorizing a new clip, Factorsynth's playback will stop.

 When launching a factorization, all connections and element edits are reset. Think about saving your project or exporting the components before launching a new factorization.

Factorizing an x-syn sound

To factorize an x-syn sound, just select the components from the purple drop-down menu and drag and drop any clip or sound file to the purple area to the upper right. After factorization is done, you'll see the new set of spectral components displayed as purple spectra to the upper right part of the interface. You are now ready to create cross-components by clicking on the purple switchboard or operation buttons.



X-SYN

fs_clarinet.wav

If you want to re-factorize the same x-syn sound, you just need to click on the purple factorize button.




 Factorization of the x-syn sound is only possible if a main sound has been previously factorized.

Progress bar










The progress bar shows the computation progress of the factorization and file export processes. During computation, a cancel button will show up on the right end of the progress bar. You can cancel computation anytime by clicking on the cancel button.


 Factorization can be a CPU and memory-demanding process. For short short clips (less than 30 seconds) and few components (less than 8), factorization will only take a few seconds. For longer clips with more components, factorization time might be long, and memory can go up to several GBs! Factorsynth does not impose any limit on the clip length, it is up to you to check memory usage and decide how far you can go! If something is too slow or memory-demanding, remember that you can always cancel factorization once started.

Switchboard operations

As an alternative to creating components by clicking on individual switchboard buttons, the main and x-syn operation buttons automatically create/remove several components at once:

	Diagonal	Add all the diagonal components. For the main sound, this means: reconstruct the original sound. For the x-syn sound, it creates a “frequency-coherent” hybrid sound: since the spectral shapes are sorted by brightness, the brightest (resp. darkest) spectral shapes of the x-syn sound will get associated with the brightest (resp. darkest) shapes of the main sound.
	Scramble	Create random connections, but keeping only one connection per spectral shape.
	Random	Create random connections, allowing connecting several temporal elements to a single spectral shape, and vice-versa.
	Clear	Clear all connections.
	All	Activate all connections.
	Clear diagonal	Clear the diagonal connections. For the main sound, this means: remove all original components.
	Add/remove one	Randomly add/remove single components.

Solo buttons




The solo buttons activate “solo mode” for either the main or the x-syn sound. Solo mode allows listening to individual components. When in solo mode, clicking on individual components on either switchboard activates that component and clears all the others. When exiting solo mode, the switchboard connections that were active before entering solo mode are restored.

Main transport

To playback Factorsynth’s processed output, use the play/stop button on the main transport panel. The panel contains two other controls:



- **Loop:** playback will loop around the output waveform.
- **Host:** in “host” mode, Factorsynth’s output will synchronize with the tempo and beats of the Live set, otherwise it will play back at the clip’s original speed. This is equivalent to Live’s global “warp” setting for clips. When clicking on play/stop in “host” mode while the Live set is playing, playback/stopping will happen at the next measure start (the button blinks meanwhile).


- **:2** and ***2** buttons (only available in host mode): play at double or half the clip’s detected tempo
- **Beats**  : the beat display window shows the number of beats detected from the clip. When in host mode, you can click and edit the number of beats to change how the clip is synchronized with the Live set.

Playback input sounds




The small play buttons above the clip name fields play back the original loaded clips at the original speed.

Input waveform (main)



The waveform of the loaded main sound, previous to modifications, is displayed in blue.

Output waveform



The current output waveform (the sum of all main and x-syn components) is displayed in yellow, with the left channel waveform displayed in bright yellow and the right channel in dark yellow.

You can select a playback region by dragging over the output waveform. If in loop mode, playback will loop inside that region. To go back to playing back the whole waveform, click on the **ALL** button above the waveform.

To quantize a playback region to subdivisions of the beats, as determined by the quantization drop-down menu, click on the **SNAP** button before selecting the region. Vertical bars corresponding to the beats will be displayed. Note that, in order for this to align with the Live set, the loop has to contain a whole number of beats, and playback must be in "host" mode.

Element mixing and editing






Apart from adding or removing components as a whole, Factorsynth allows several editing and mixing operations for individual temporal or spectral shapes.

Individual temporal shape controls




To the left of each temporal shape is a panel of mixing and editing operations that apply to all components connected to that temporal shape. In other words, all changes to the panel affect the current whole row of main and x-syn components on the switchboards.


These controls are the following:

	Level	The level dial controls the output level of all components connected to that row. Double-click on the dial to go back to the original (center) setting.
	Pan	The pan slider controls the stereo panning for all components connected to that row. Double-click on the slider to go back to the center position.
	Shift	The shift buttons displace the current temporal shape to the left/right. Use this to create rhythmic variations of loops. The amount of shifting is determined by the global "quantization" setting (see below). Shifting is circular: when shifting to the right, the vanishing right end of the shape comes back on the left, and vice-versa.
	Threshold	The threshold button removes low-energy residuals from the current temporal shape. Use this to clean leftover noises between sound events, such as between drum hits.
	Reset	The reset button restores the current temporal shape to its original state right after factorization.



Global time shape controls



The controls on the global time shape panel affects all time shapes at once. The **threshold** and **reset** buttons have the same effect as described above, but applied to all time shapes. Processing might take a while.


The **quantization** (quant.) setting determines the allowed amount of shifting when using the individual shifting buttons . It is given in fractions of a beat. There is also a "none" setting, in which case the temporal shapes can be shifted arbitrarily.

Spectral shape controls

Spectral shapes have only a level dial  and a reset button . In this case, this will affect all the components connected to the current *column* of either the main or the x-syn switchboard.


Free shape editing

Finally, temporal and spectral shapes can be edited by **drawing** on them. The changes will be effective when you release the mouse button after drawing. Note that the reset button also restores all the drawn edits, and that saving of drawn edits with the Live set is currently not supported.




Analysis panel

The analysis panel contains four parameters that, together with the number of components, control the factorization algorithm. Note that, when you modify one of them, it will be applied on the *next* factorization you will perform.



- **Resolution** (res.): sets the time resolution (in milliseconds) of the analysis window. Possible values: 5-180. Default: 80. Low values mean a high temporal resolution (e.g., attacks and impulses will sound cleaner), but at the cost of a lower frequency resolution (e.g. it will be harder to distinguish individual harmonics).
- **Hop**: sets the hop size (given in % of overlap) of the analysis window. Possible values: 0.25 and 0.125. Default: 0.25. A lower value means a higher temporal resolution without changing the frequency resolution, but at the cost of higher CPU and memory demands.
- **Algorithm**: it is possible to choose between two variations of the NMF decomposition algorithm. Algorithm B has a higher tendency to keep transients and stable parts of notes together in the same component.
- **Iterations**: Possible values: 5 to 200. Choosing a high number of iterations will often result in a better separation of the factorized elements, but at the cost of longer computation times.

Clear panel


The clear panel allows clearing all temporal and spectral shapes and the output components. In contrast to the "clear" operation buttons  , this not only clears the switchboards, but also the temporal and spectral shapes, as well as the sound output.




- **Clear x-syn**: clears x-syn's spectral shapes, its switchboard, and the corresponding contribution of the x-syn sound from the output.
- **Clear all**: clears all temporal (main) and spectral (main and x-syn) shapes, both switchboards, and the sound output.

Export panel


The export panel controls the file exporting options of Factorsynth. All exports are saved as 16-bit WAV files with the same sampling rate than the main sound.



- **Export selected components**: a WAV file is created for each main and x-syn component that is currently selected on the switchboards. Clicking on the button opens a dialog window to select the saving location. Factorsynth will create a folder named **fsComponents** at this location, which will contain the WAV files with the following naming convention: **m_i_j.wav** for the main components and **x_i_j.wav** for the x-syn components, where **i** and **j** are the row and column coordinates of the components on the switchboards, (1,1) being the upper-left component.

 Tip: once the component WAV files have been exported, you can import them back into the Live project as multiple new tracks in arrangement view by selecting them all on the Finder (or Explorer) or on Live's file browser, and dragging them into the project while pressing the command key (⌘) on Mac or the CTRL key on Windows.

- **Export all components**: a WAV file is created for all main and x-syn components, whether they are selected or not, using the same naming convention as above.


 Be careful when exporting all components. Even a limited number of factorization components can produce a huge number of output files. For instance, having 8 main components and 8 x-syn components would produce $8*8*2 = 128$ WAV files! You can always abort exporting via the cancel button at the right of the progress bar.

- **Bounce to file:** a WAV file is created containing the current output sound (the waveform displayed in yellow). The default output file name is `fs_out.wav`

Mixer panel

The mixer panel controls the main output volume of the device.

- **Main volume:** controls the volume of the part of the output produced by the main components.
- **X-syn volume:** controls the volume of the part of the output produced by the x-syn components.
- **Mute buttons (main and x-syn):** mute all the main resp. x-syn components. As opposed to clearing the switchboards, muting eliminates the output but does not clear the selected components.
- **Dry/wet:** mix of the original, unprocessed main sound, with Factorsynth's output.
- **Gain:** the output gain of the device. Double-click on the dial to get back to original (0dB) gain.



Working with presets

You can use Live's integrated preset saving system (clicking on the floppy disk icon on the upper right of the device window in the audio effects area). Saving a Factorsynth preset stores the switchboard connections, the dials and panning boxes, the quantization setting and the analysis parameters. Saving of free (drawing) edits to the temporal and spectral shapes and of temporal shifts is not supported.

The same applies when saving and re-loading a Live Set that includes a Factorsynth device.

When re-loading a preset, factorizations need to be performed again. They will be automatically launched right after loading.

Automation and MIDI-mapping

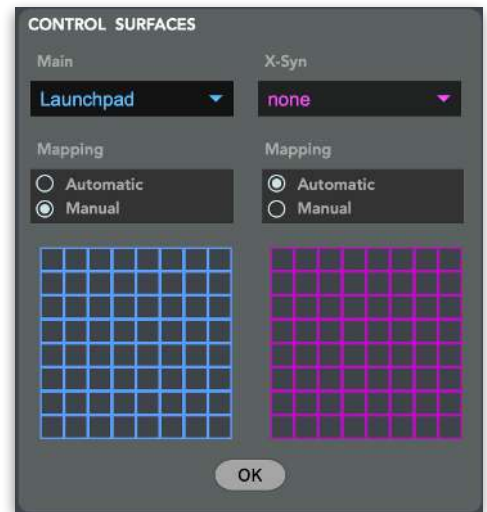
The following Factorsynth parameters are controllable by Live's automation envelopes and MIDI and key mappings: first 8x8 individual buttons on the switchboards, factorization buttons, matrix operation buttons, solo buttons, number of components, all levels and pan sliders, shifting buttons, transport controls and quantization.

The following are not controllable: analysis parameters, clear panel, export panel, "snap" and "all" buttons on the output waveform display, reset and threshold buttons and individual buttons on the switchboards corresponding to component numbers higher than 8.

- **Automatic MIDI mapping of the switchboards**

Some surface controllers (e.g. Launchpad, Push...) can be used to control the first 8x8 shapes of the main or x-syn switchboards without the need for any manual MIDI mapping:

1. Click on the **control** button on the upper part of the main interface to access the surface control configuration window.
2. From the dropdown menus, select which control surface is to be assigned to which switchboard.
3. Make sure that the "Mapping" mode is set to "Automatic".



- ⚠ Some control surfaces are not recognized by their names and might be listed as "GenericScript".
- ⚠ If automatic mapping is not working, try switching between different mode pages if your control surface supports them. For example, Launchpad devices will work with Factorsynth on the "User 2" page.

- **Manual MIDI mapping of the switchboards**

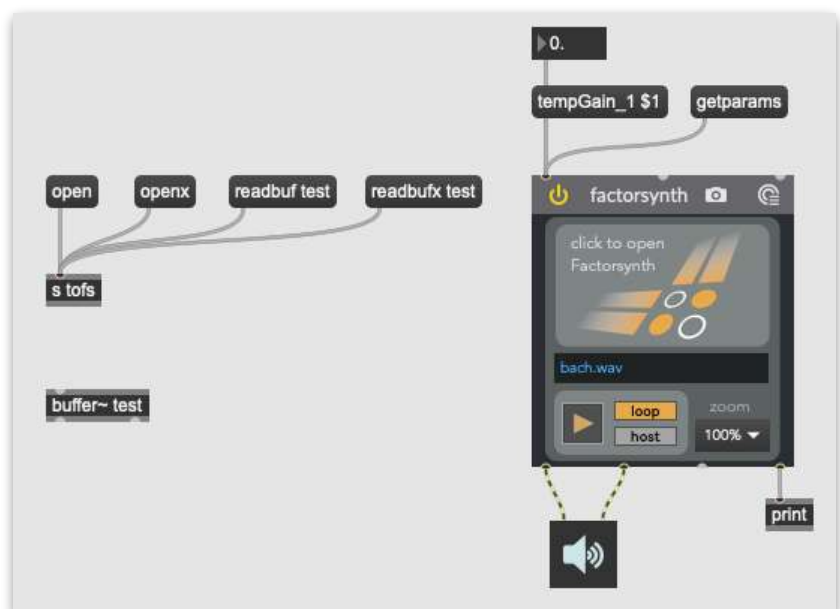
If your control surface does not work in automatic mapping mode, you can still manually assign the first 8x8 buttons of both the main and x-syn switchboards:

1. Click on the **control** button on the upper part of the main interface to access the surface control configuration window (see image above).
2. Select "Manual" mapping mode for the corresponding switchboard.
3. Click on the MIDI mapping button on the upper right part of Live's interface. You will see that the blue and purple grids are highlighted.
4. One by one, click on the button to be mapped on the grid, followed by pushing the corresponding control surface button. You should see small message windows appearing over each correctly mapped button.

Usage from Max

It is possible to load and use Factorsynth from plain Max, without Ableton Live, thanks to the `amxd~` object. To load it into a patcher, just drag and drop the `factorsynth.amxd` file onto it. You'll see a window with the Factorsynth logo. To open the main interface, click on the logo.

Connect stereo audio out to the first two outlets. Any parameter that can be automated in Live can be modified by sending a message to the left inlet containing the parameter name and value



(as shown in the figure for `tempGain_1`). To get a list of all available parameters, send a `getparams` message to the left inlet and read the results via a `print` object connected to the right outlet.

There are two ways of importing sounds into Factorsynth from Max, controlled by the following messages that need to be passed to a send object named `stofs`, as shown in the figure:

- `open` and `openx`: load a sound file for, respectively, the main and x-syn sounds.
- `readbuf` and `readbufx`, followed by the name of a `buffer~` object: load the sound stored in that `buffer~` object as, respectively, the main or x-syn sound.

Factorization how-to

Choosing the right number of components is crucial for obtaining good factorizations. Of course, this will depend on the complexity of the input sound and on the type of elements you would like to obtain. You'll see that it is mostly a trial-and-error process, but here are some guidelines that might help you:

- If your input clip contains a few simple sound events (for example, a drum loop or a sequence of a few notes), Factorsynth will likely separate them into components. For example, if you have a melody containing 7 different pitches (not counting the repeated ones), doing a factorization with 7 components will often (but not always!) result in one temporal/spectral shape pair per pitch.
- In the same scenario of a simple clip with few events, if you factorize into more components than events in the original sound, Factorsynth will start "delving" deeper to extract more underlying structures. For instance, if your clip contains 3 piano notes and you factorize into 10 components, some components will contain the noisy and impulsive parts of the attack phases, and other will contain the harmonic sustain and/or release parts of the different pitches.
- For complex input clips (loops containing several instruments, song excerpts or full songs), the decomposition is more unpredictable, but can often detect interesting elements such as a rhythmical motif of the bass, a guitar lick, an underlying pad, impulsive noises, etc.
- If you decompose a song excerpt with few components, it is likely that all the drums will end up on one or two components.
- Factorsynth cannot separate full instruments or voices, except in very simple cases. You won't be able to do things like suppressing a singing voice or guitar solo. That's the job of *source separation*, which is a harder thing to do! In that sense, Factorsynth is not a "source separator", it is a tool for sound design based on deconstructing sounds into interesting (and sometimes unexpected!) elements.

As you can see, using Factorsynth is a pretty exploratory process. After a bit of testing, you will get a sense of how it reacts to the clips in your collection.

From Factorsynth 1 to Factorsynth 2

Factorsynth 2 has introduced several important changes in terms of workflow. All of them were inspired by a lot of helpful feedback from users of version 1.x. Here is a short transition guide for those used to Factorsynth 1:



- Main sounds are now always imported by drag and drop, instead of selecting the clip on the Live set and clicking on "factorize".
- When a sound is imported, factorization is instantly launched. You only need to click on the "factorize" buttons if you wish to re-factorize an already-loaded sound with new parameters.

- Factorsynth 2 has its own transport controls (play/stop/loop), instead of relying on the position and looping settings of the original clip on the Live set for playback. This allows a much more robust synchronization with the Live set when using Factorsynth in "host" mode, as well as keeping up with real-time changes of the set tempo. If you need to start playback at a specific time, you can use an automation curve to drive the play button.
- There is no longer the need for real-time "listening" of warped clips. They are now instantly factorized, just like unwarped clips. The same goes for encrypted samples from some Ableton packs.
- Version 2 allows to select between two different factorization algorithms. Algorithm "A" is the exact same algorithm that was present in version 1.
- A new gain dial allows to boost the output somewhat, if needed.
- Apart from these workflow changes, the new features in version 2 are: temporal shifting, panning, thresholding, all-on operation, add/remove random components and playback region selection.

Support

For assistance, bug reports and any suggestions, please contact software@jjburred.com

If you're sending in a bug report, please include the following information:

- Your Factorsynth version
- Your OS version
- Your Ableton Live version
- Your Max (Max For Live) version (this can be found by opening the Max editor by clicking the  button on the device window)
- Any possible error messages displayed in the Max console, displayed by clicking on the  button on Factorsynth
- If possible, the audio clip that produced the issue