

Coagula

Concatenative player for one long sample.

Load one long sample (the target is about an hour). It is analysed once into four segmentations plus per-segment descriptors, then played back by descriptor-space retrieval under live modulation. You steer a target point in a scatter plot of the sample's segments; the nearest segments get played.

The engine runs in Node for Max on worker threads (decode, analysis, sidecar cache, k-nearest-neighbour retrieval). Playback is a 32-voice gen~ segment player. A one-hour stereo WAV analyses in about 8 s on the dev Mac and reloads from its cache in well under a second.

ARCHITECTURE

```
sample -> Node engine: decode, 4 segmentations, descriptors, sidecar cache
          | target X/Y + Jitter + k -> kNN retrieval
          v
MIDI / Chain / Clock -> gen~ player: 32 voices, equal-power crossfades, peak limiter
```

QUICK START

- 1 Load Coagula on a MIDI track. Drop an audio file on the Source well, or use **Load**. Analysis progress shows in the well.
- 2 Choose **Trigger**: Chain plays segment after segment, Clock plays one per tick (set Rate, or Sync), MIDI plays on note-on.
- 3 Press **Run** (Chain and Clock), then click or drag in the scatter plot to move **Target X / Y**.
- 4 Move **Morph** between Assembly (recognisable reordered pieces) and Texture (dense, short, jittered grains).
- 5 In Pitch axis mode, turn on **Key Trk** to have MIDI notes pick the pitch target; add **Retune** to transpose the segment to match.

CONTROLS

SOURCE AND SPACE

CONTROL	RANGE	WHAT IT DOES
Sample / Load	-	Drop target and Load button. The path is stored with the Set; Collect All and Save does not collect the sample.
Variant	Onset / Hyb 400 / Hyb 200 / Fixed	Which of the four segmentations plays. Switching is a hard cut; the playing segment finishes first.
Axis	Pitch / Noise	What the X axis measures: pitch, or noisiness (spectral flatness).
Target X / Y	0-100 %	The target in descriptor space (X pitch or noise, Y spectral centroid), normalised to the variant's range. Dragging the plot writes these.

PLAY

Morph / Link	0-100 %	Macro from Assembly (0) to Texture (100) over k, Jitter, Crossfade, Continuity and Max Len. While Link is on the bundled dials follow Morph; off, they are independent.
Silence / Match	0-100 %	Chance of picking a pause segment. Match resets to the sample's own pause ratio.
Unvoiced	0-100 %	Chance of an unvoiced segment (Pitch axis only).

GRAIN

k	1-32	Pick uniformly from the k nearest segments.
Jitter	0-50 %	Random offset of the target, as a fraction of the axis range.
X-Fade	2-200 ms	Crossfade between segments (clamped to half the segment length). Default 8.
Contin.	0-100 %	Chance of playing the segment that follows the last one in source order. Default 60.
Max Len	20-1000 ms, top = Off	Truncates long segments at playback.
Guard	0-16	Excludes the last N segments, so a static target does not loop one grain. Default 4.
Sil Thr	-70 to -20 dB	Level below which a segment counts as silence. Default -50.
Transp	-24 to +24 st	Playback transposition.

TRIGGER AND OUTPUT

Trigger	Chain / Clock / MIDI	Chain starts each segment at (end - crossfade). Clock consumes one per tick. MIDI commits on note-on.
Run, Sync, Rate, Division	0.5-100 Hz	Clock runs while Run is on or a note is held; with Sync it also needs the transport running and uses Division. Default 8 Hz.
Key Trk, Retune	toggles	Key Trk makes the MIDI note an absolute pitch target. Retune resamples the segment to match (clamped to +/-12 st).
Gain	-70 to +6 dB	Output fader with meter.

NOTES

Analysis page

The **Analysis** button reveals the settings used when a sample is analysed: Frame (1024 / 2048 / 4096), Hop (128 / 256 / 512), Onset threshold, Pause Gate, Min Seg, Fixed Chop, the two Hybrid split lengths, and Pitch Lo / Hi. Changing them re-analyses the sample. They are stored with the Set but are not automatable.

Selection order

Each pick runs: Silence roll, Continuity roll, Unvoiced roll (Pitch axis only), then the main nearest-neighbour query at the target plus Jitter, with the Repeat Guard excluding recent segments. An empty bucket falls through to the next step; it never stalls.

Limits

- macOS only as built: the engine is extracted to /tmp at load, and non-WAV/AIFF files are converted once with afconvert.
- No MuBu: its externals cannot be embedded in a single frozen device, so analysis, persistence and retrieval run in Node.

- Pitch is YIN on a decimated copy; onsets are log spectral flux with an adaptive threshold. Segment descriptors are tuned for speech and solo saxophone.

STATUS

Checked in Live 12.4.6: the engine starts from the frozen device; a speech file and a one-hour stereo WAV analyse, cache and reload; Chain, Clock and MIDI (with key tracking and Retune) play; Morph drives the dials; scatter clicks set Target X/Y. Not yet checked: Set save/reopen restoring the path, read-only folder fallback, Sync against the transport, and anything by ear beyond the meters.