

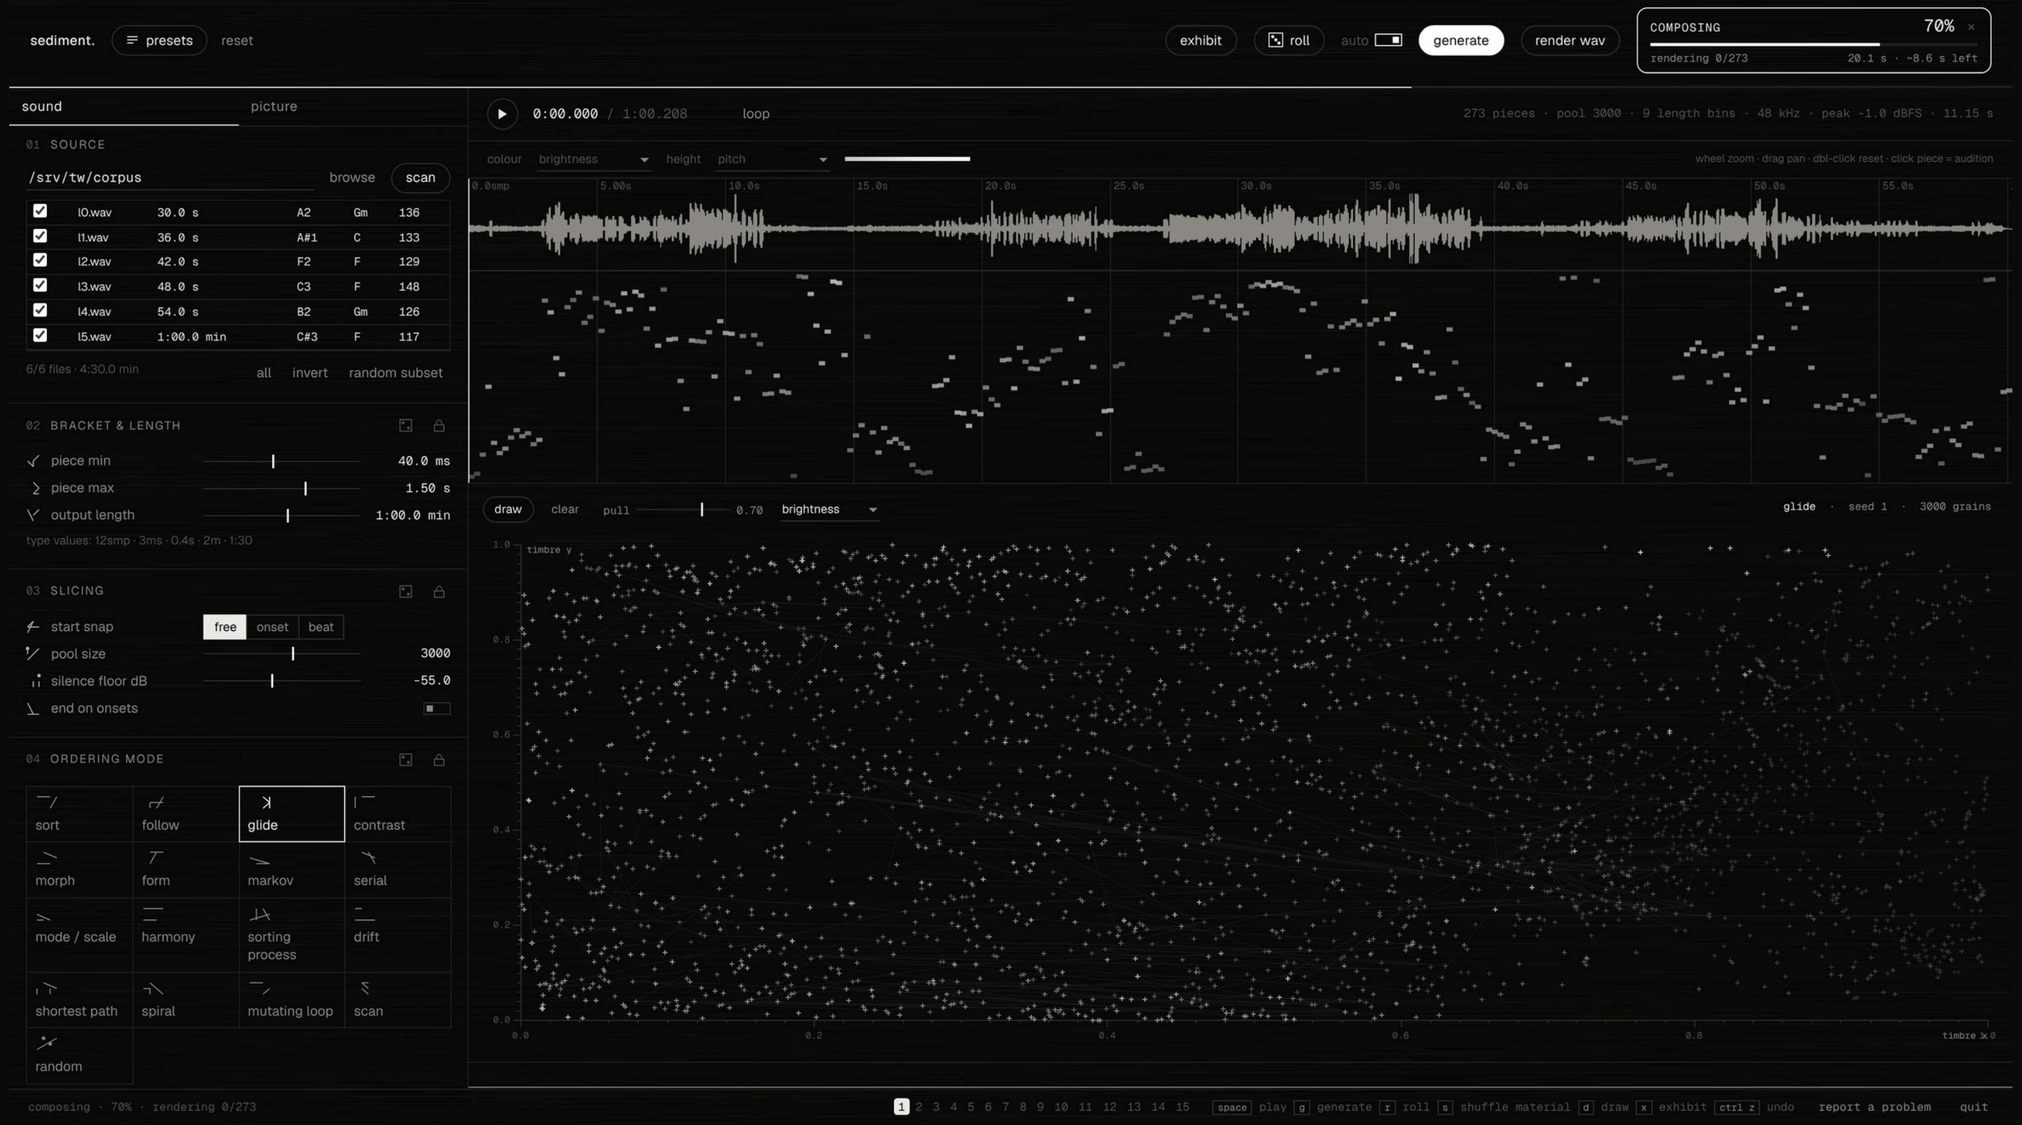
Thomas Wagensommerer

OPEN SOURCE APPLICATIONS

COMMUNAL TOOLS 2026

INDEX

T01	sediment	Sound
T02	sediment·live	Sound
T03	sediment·fx	Sound
T04	sediment·pod	Sound
T05	ligature	Picture-bound sound
T06	sortcut	Picture-bound sound
T07	scorecut	Picture-bound sound
T08	datamosh station	Moving image
T09	scanline	Moving image
T10	framegrab	Moving image
T11	AVjoin	Moving image
T12	dubtrans	Voice and language
T13	subtrans	Voice and language
T14	voicebox	Voice and language
T15	TerraForge	Space, terrain and body
T16	meanland	Space, terrain and body
T17	GNM head pipeline	Space, terrain and body
T18	AutoRagdoll	Space, terrain and body
T19	Resonode	Performance control



sediment

2026

SOUND

01 · sound · glide across timbre space

T01

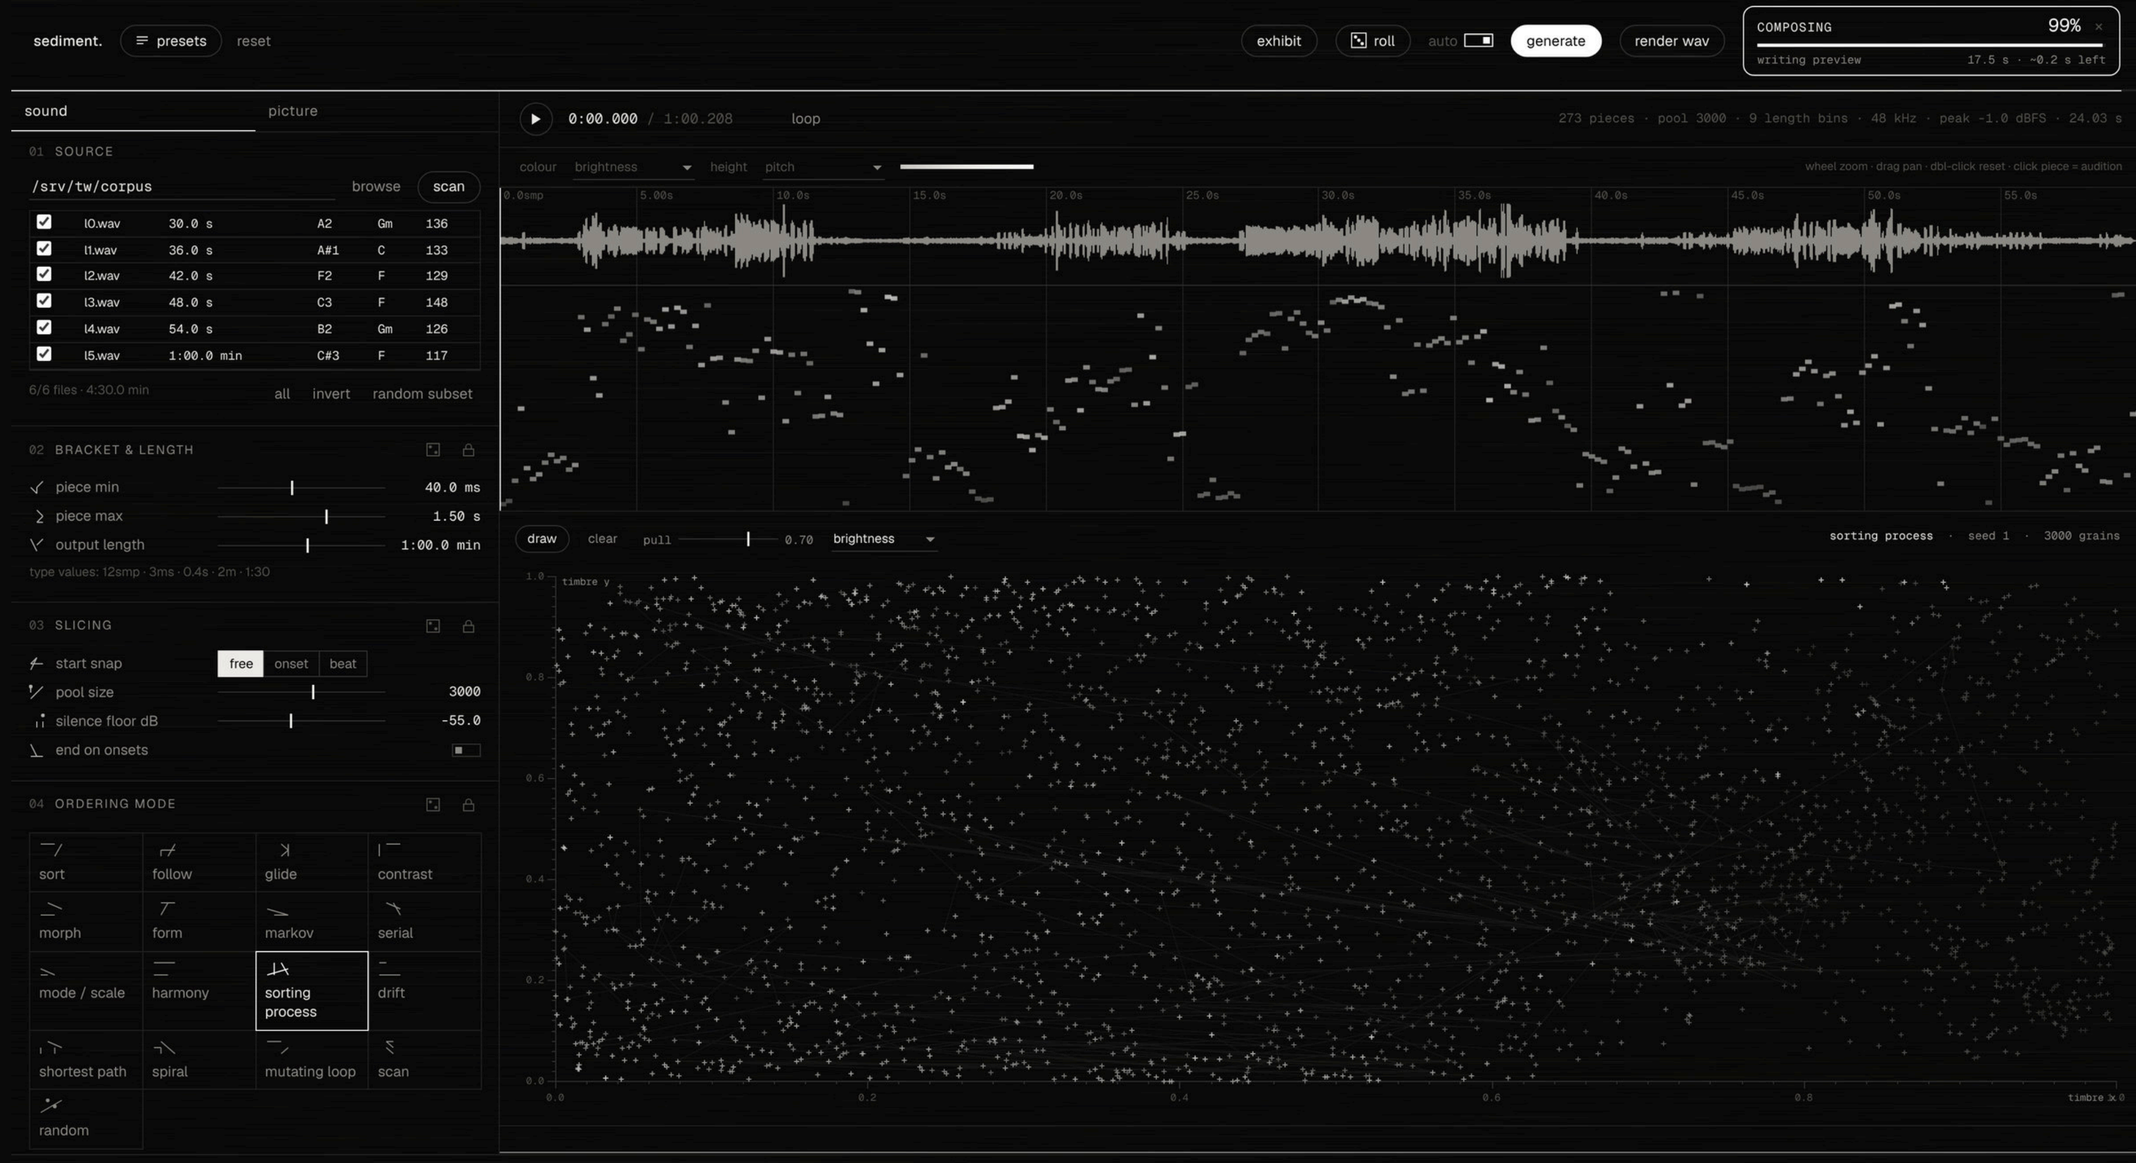
T01

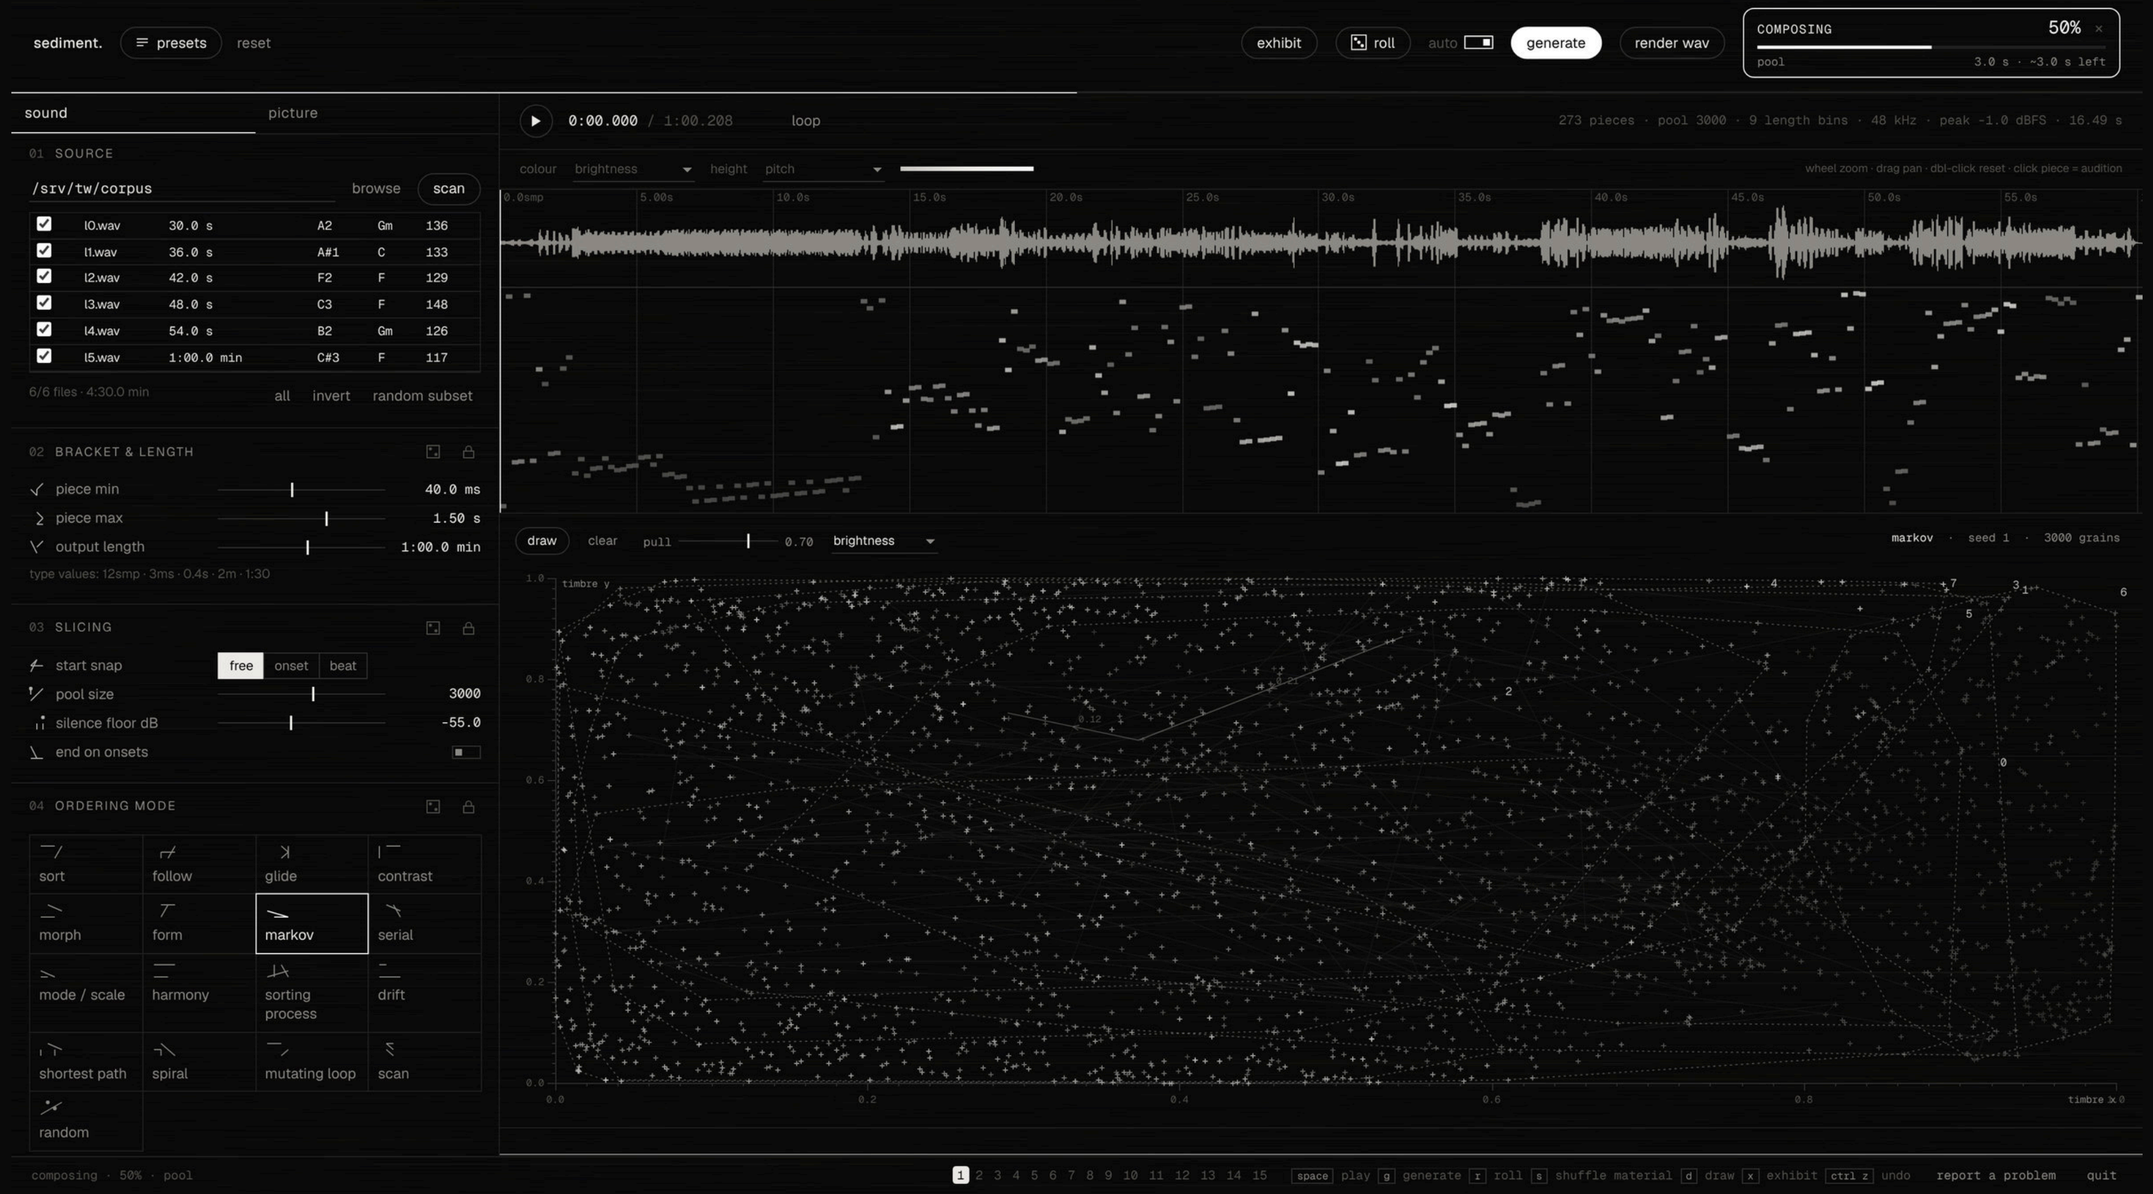
sediment

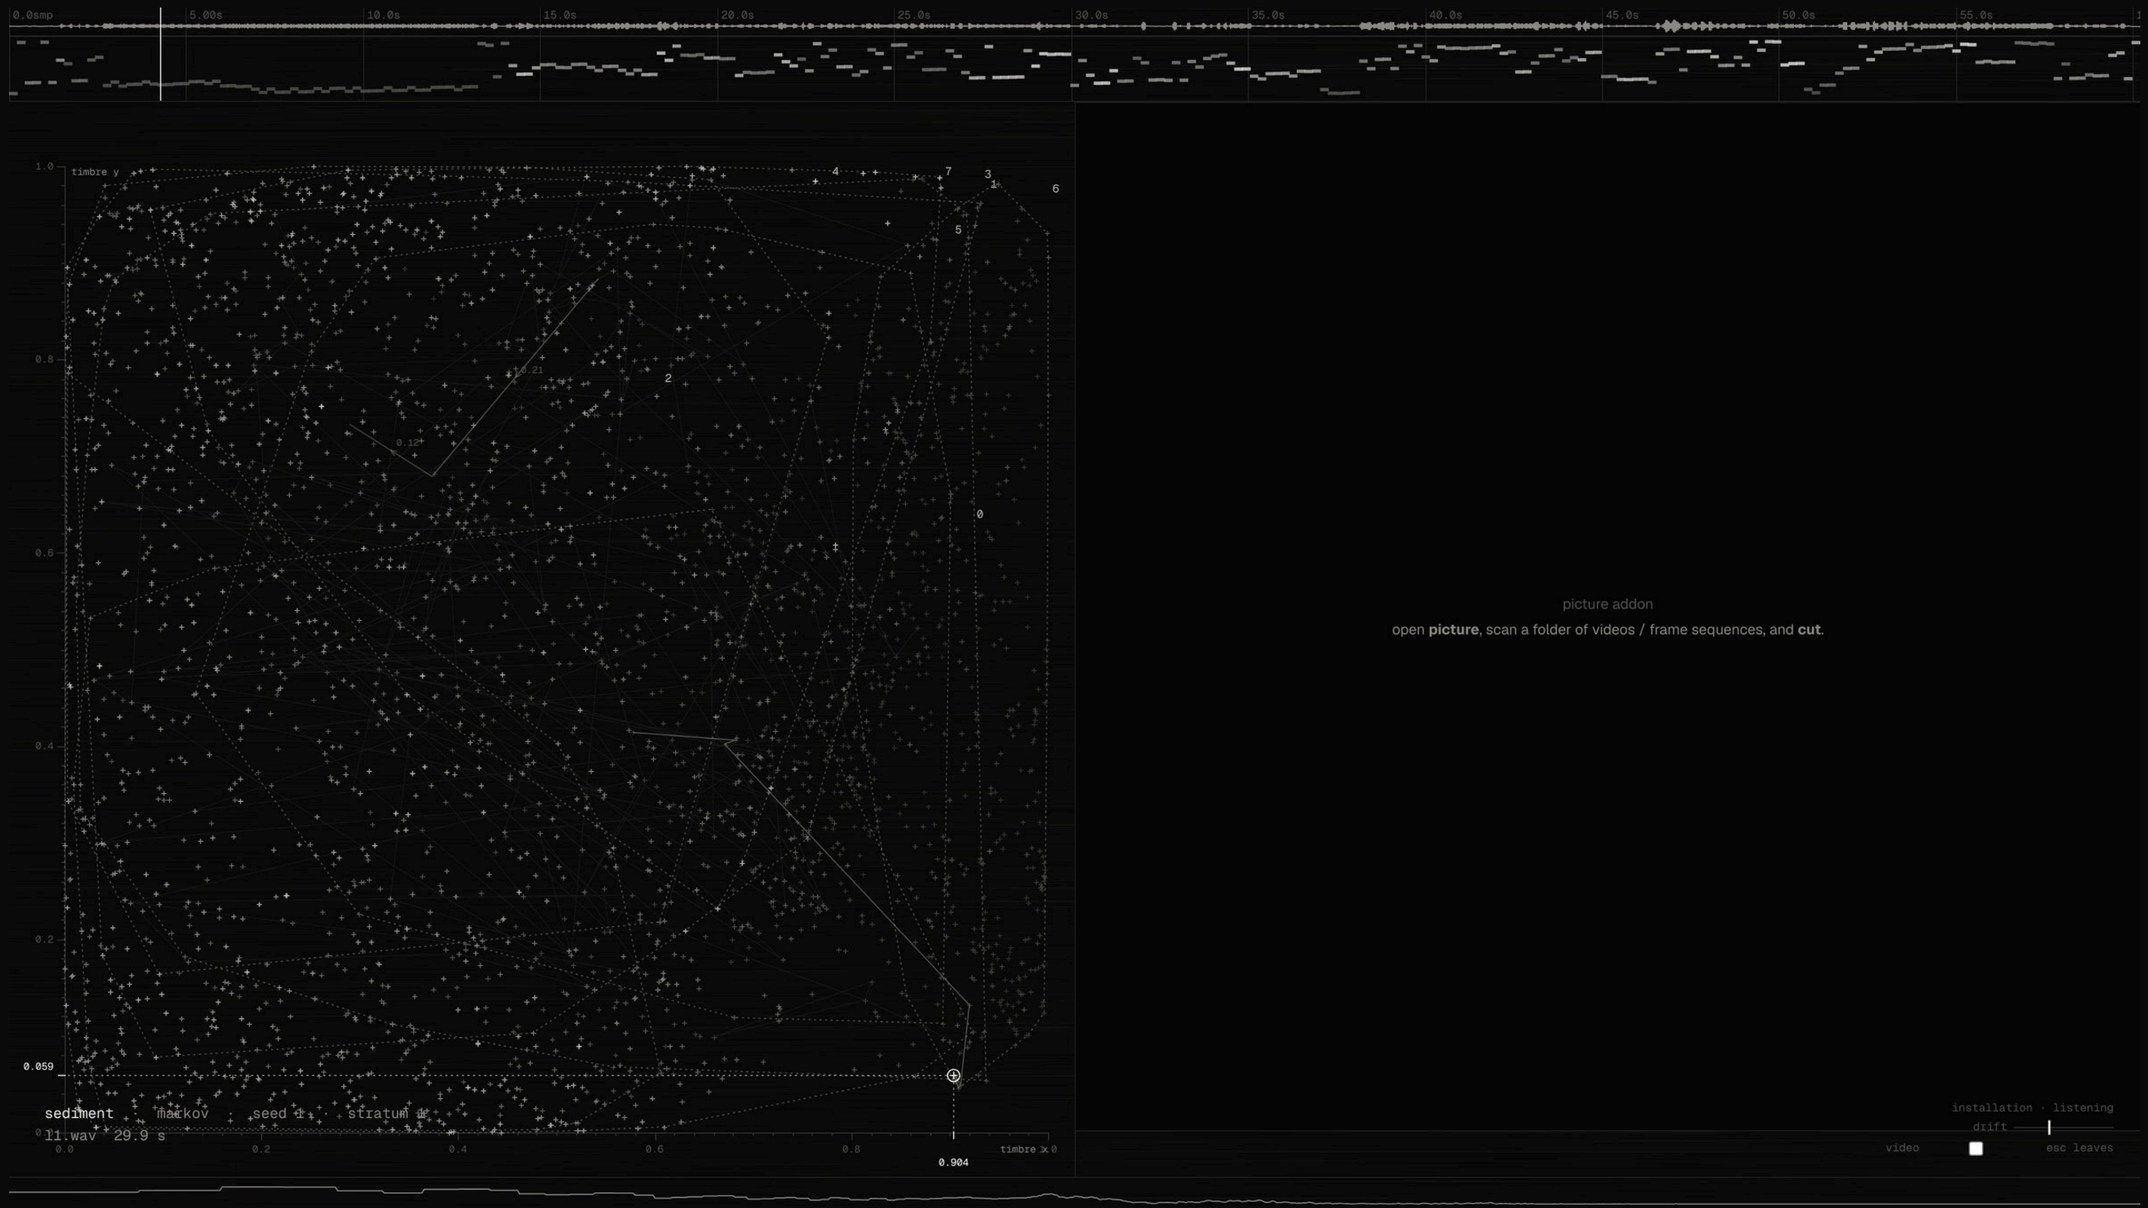
2026

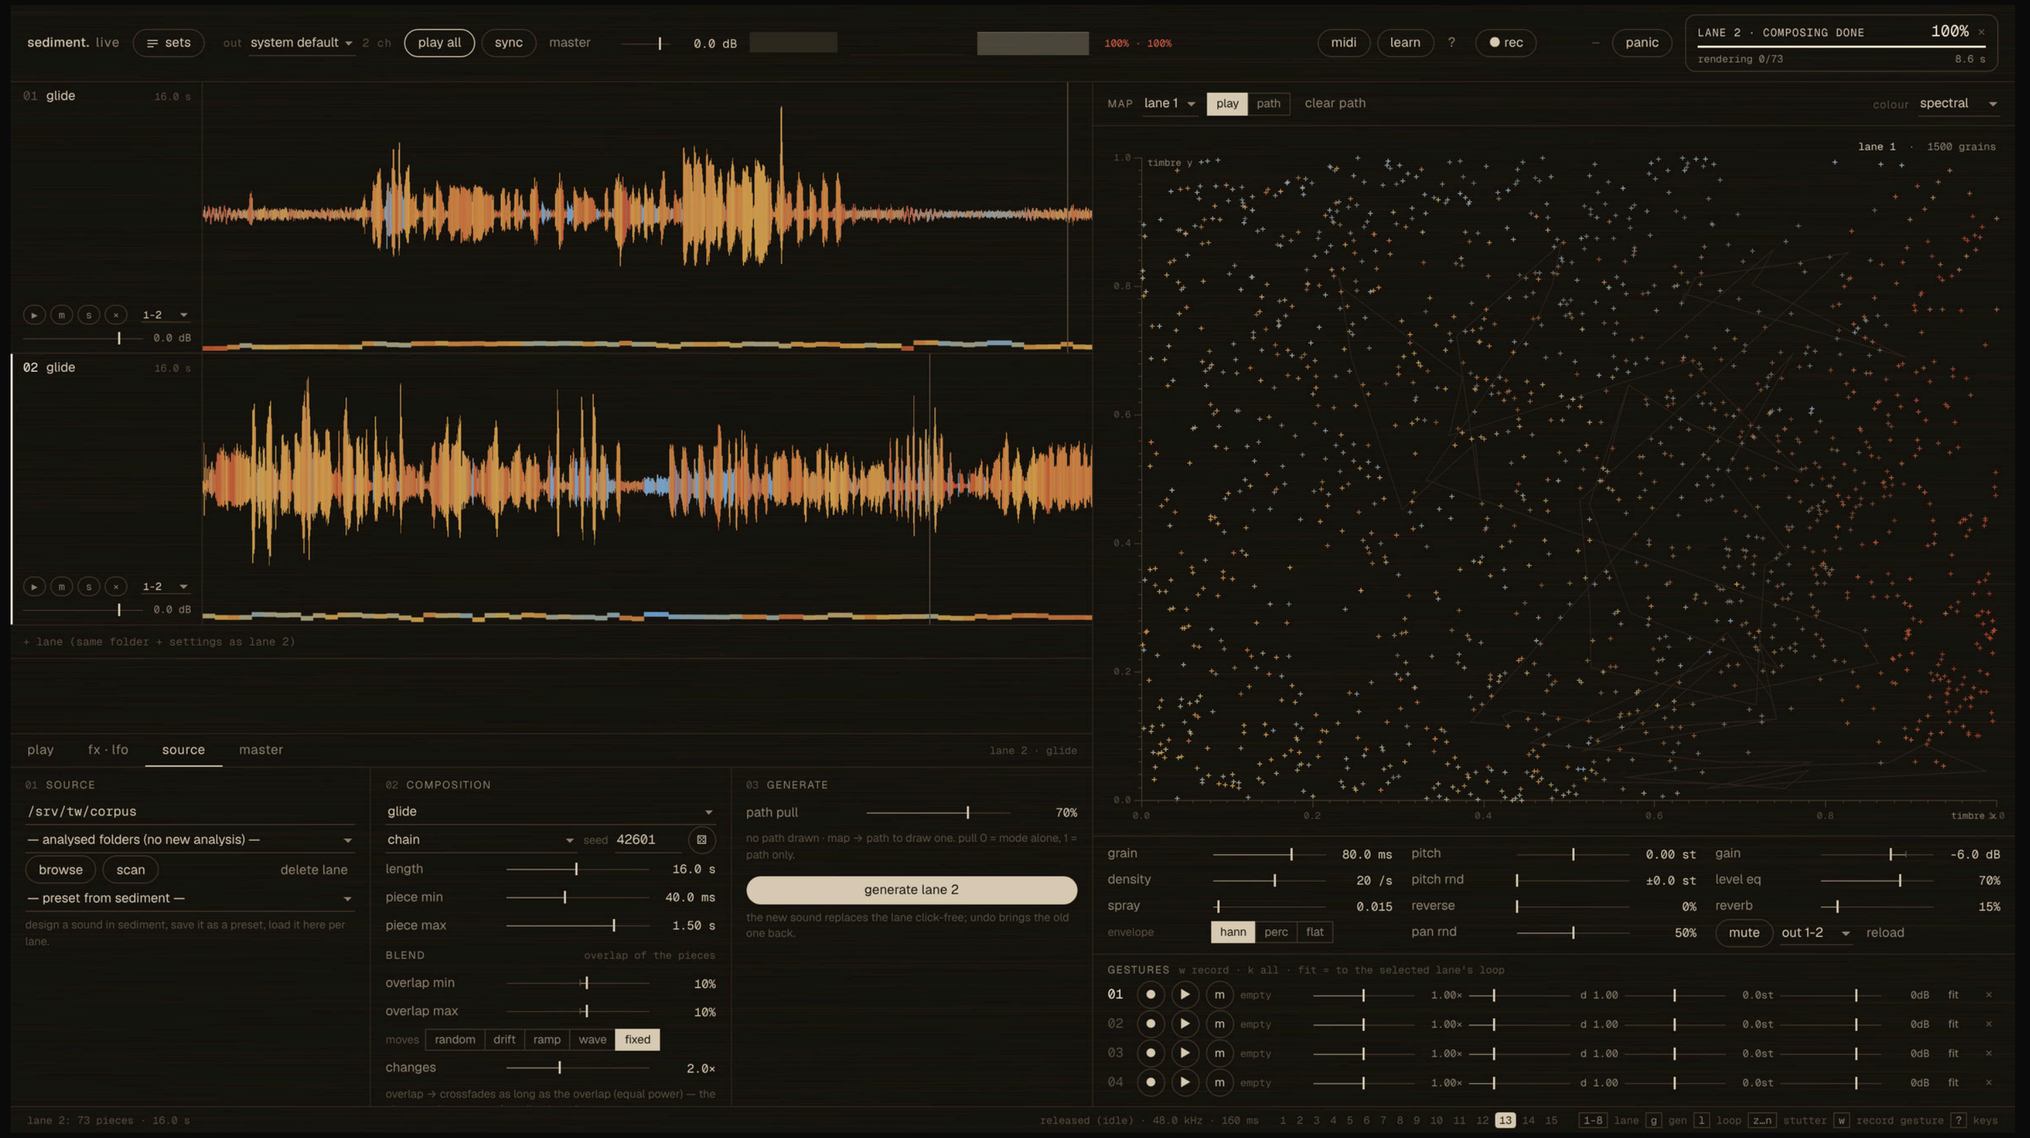
FIELD	Sound — corpus-based composition and live processing
VERSION	1.0
TECHNOLOGY	Python (numpy/scipy, aiohttp), browser UI, ffmpeg; Windows/macOS self-installer
SOURCE	sediment-1.0.zip

A composition environment that treats a folder of recordings as a corpus. Every file is analysed into frame-level descriptors (loudness, onset density, pitch and pitchedness, spectral shape, chroma and key, stereo width, an MFCC-derived timbre space); pieces of arbitrary length are then re-ordered according to explicit ordering principles — sorting, gliding, contrast, morphing, formal clustering, Markov grammars learned from the sources, serial rows, modal and harmonic retuning, audible sorting algorithms — and rendered as high-resolution audio with a full event log. A picture add-on cuts frame-exact video to the resulting sound. The tool makes the ordering itself the compositional material: structure becomes audible as process.


sediment.
presets
reset
exhibit
roll
auto
generate
render wav
COMPOSING
99%
writing preview
17.5 s · ~0.2 s left
sound
picture
01 SOURCE
/srv/tw/corpus
browse
scan
l0.wav 30.0 s A2 Gm 136
l1.wav 36.0 s A#1 C 133
l2.wav 42.0 s F2 F 129
l3.wav 48.0 s C3 F 148
l4.wav 54.0 s B2 Gm 126
l5.wav 1:00.0 min C#3 F 117
6/6 files · 4:30.0 min
all
invert
random subset
02 BRACKET & LENGTH
piece min 40.0 ms
piece max 1.50 s
output length 1:00.0 min
type values: 12smp · 3ms · 0.4s · 2m · 1:30
03 SLICING
start snap
free
onset
beat
pool size 3000
silence floor dB -55.0
end on onsets
04 ORDERING MODE
sort
follow
glide
contrast
morph
form
markov
serial
mode / scale
harmony
sorting process
drift
shortest path
spiral
mutating loop
scan
random
0:00.000 / 1:00.208
loop
273 pieces · pool 3000 · 9 length bins · 48 kHz · peak -1.0 dBFS · 24.03 s
colour brightness
height pitch
wheel zoom · drag pan · dbl-click reset · click piece = audition
0.0smp
5.00s
10.0s
15.0s
20.0s
25.0s
30.0s
35.0s
40.0s
45.0s
50.0s
55.0s
draw
clear
pull
0.70
brightness
sorting process · seed 1 · 3000 grains
timbre y
timbre x
composing · 99% · writing preview
space play g generate r roll s shuffle material d draw x exhibit ctrl z undo
report a problem
quit


sediment.
presets
reset
exhibit
roll
auto
generate
render wav
COMPOSING
50%
pool
3.0 s · ~3.0 s left
sound
picture
01 SOURCE
/srv/tw/corpus
browse
scan
l0.wav 30.0 s A2 Gm 136
l1.wav 36.0 s A#1 C 133
l2.wav 42.0 s F2 F 129
l3.wav 48.0 s C3 F 148
l4.wav 54.0 s B2 Gm 126
l5.wav 1:00.0 min C#3 F 117
6/6 files · 4:30.0 min
all
invert
random subset
02 BRACKET & LENGTH
piece min 40.0 ms
piece max 1.50 s
output length 1:00.0 min
type values: 12smp · 3ms · 0.4s · 2m · 1:30
03 SLICING
start snap
free
onset
beat
pool size 3000
silence floor dB -55.0
end on onsets
04 ORDERING MODE
sort
follow
glide
contrast
morph
form
markov
serial
mode / scale
harmony
sorting process
drift
shortest path
spiral
mutating loop
scan
random
0:00.000 / 1:00.208
loop
273 pieces · pool 3000 · 9 length bins · 48 kHz · peak -1.0 dBFS · 16.49 s
colour
brightness
height
pitch
wheel zoom · drag pan · dbl-click reset · click piece = audition
0.0smp
5.00s
10.0s
15.0s
20.0s
25.0s
30.0s
35.0s
40.0s
45.0s
50.0s
55.0s
draw
clear
pull
0.70
brightness
markov · seed 1 · 3000 grains
timbre y
timbre x
composing · 50% · pool
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
space play
g generate
r roll
s shuffle material
d draw
x exhibit
ctrl z undo
report a problem
quit


0.0smp
5.00s
10.0s
15.0s
20.0s
25.0s
30.0s
35.0s
40.0s
45.0s
50.0s
55.0s
timbre y
1.0
0.8
0.6
0.4
0.2
0.059
0.0
0.2
0.4
0.6
0.8
0.904
timbre x
sediment
markov
seed 1
stratum 4
11.wav · 29.9 s
picture addon
open picture, scan a folder of videos / frame sequences, and cut.
installation · listening
drift
video
esc leaves


sediment. live
sets
out system default 2 ch
play all
sync
master
0.0 dB
100% · 100%
midi
learn
?
rec
panic
LANE 2 · COMPOSING DONE
100%
rendering 0/73
8.6 s
01 glide
16.0 s
1-2
0.0 dB
02 glide
16.0 s
1-2
0.0 dB
+ lane (same folder + settings as lane 2)
MAP
lane 1
play
path
clear path
colour
spectral
lane 1 · 1500 grains
timbre y
timbre x
play
fx · lfo
source
master
lane 2 · glide
01 SOURCE
/srv/tw/corpus
— analysed folders (no new analysis) —
browse
scan
delete lane
— preset from sediment —
design a sound in sediment, save it as a preset, load it here per lane.
02 COMPOSITION
glide
chain
seed 42601
length 16.0 s
piece min 40.0 ms
piece max 1.50 s
BLEND
overlap of the pieces
overlap min 10%
overlap max 10%
moves random drift ramp wave fixed
changes 2.0×
overlap → crossfades as long as the overlap (equal power) — the
03 GENERATE
path pull 70%
no path drawn · map → path to draw one. pull 0 = mode alone, 1 = path only.
generate lane 2
the new sound replaces the lane click-free; undo brings the old one back.
grain 80.0 ms
density 20 /s
spray 0.015
envelope hann perc flat
pitch 0.00 st
pitch rnd ±0.0 st
reverse 0%
pan rnd 50%
gain -6.0 dB
level eq 70%
reverb 15%
mute
out 1-2
reload
GESTURES w record · k all · fit = to the selected lane's loop
01 empty 1.00× d 1.00 0.0st 0dB fit
02 empty 1.00× d 1.00 0.0st 0dB fit
03 empty 1.00× d 1.00 0.0st 0dB fit
04 empty 1.00× d 1.00 0.0st 0dB fit
lane 2: 73 pieces · 16.0 s
released (idle) · 48.0 kHz · 160 ms
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
1-8 lane g gen l loop z…n stutter w record gesture ? keys

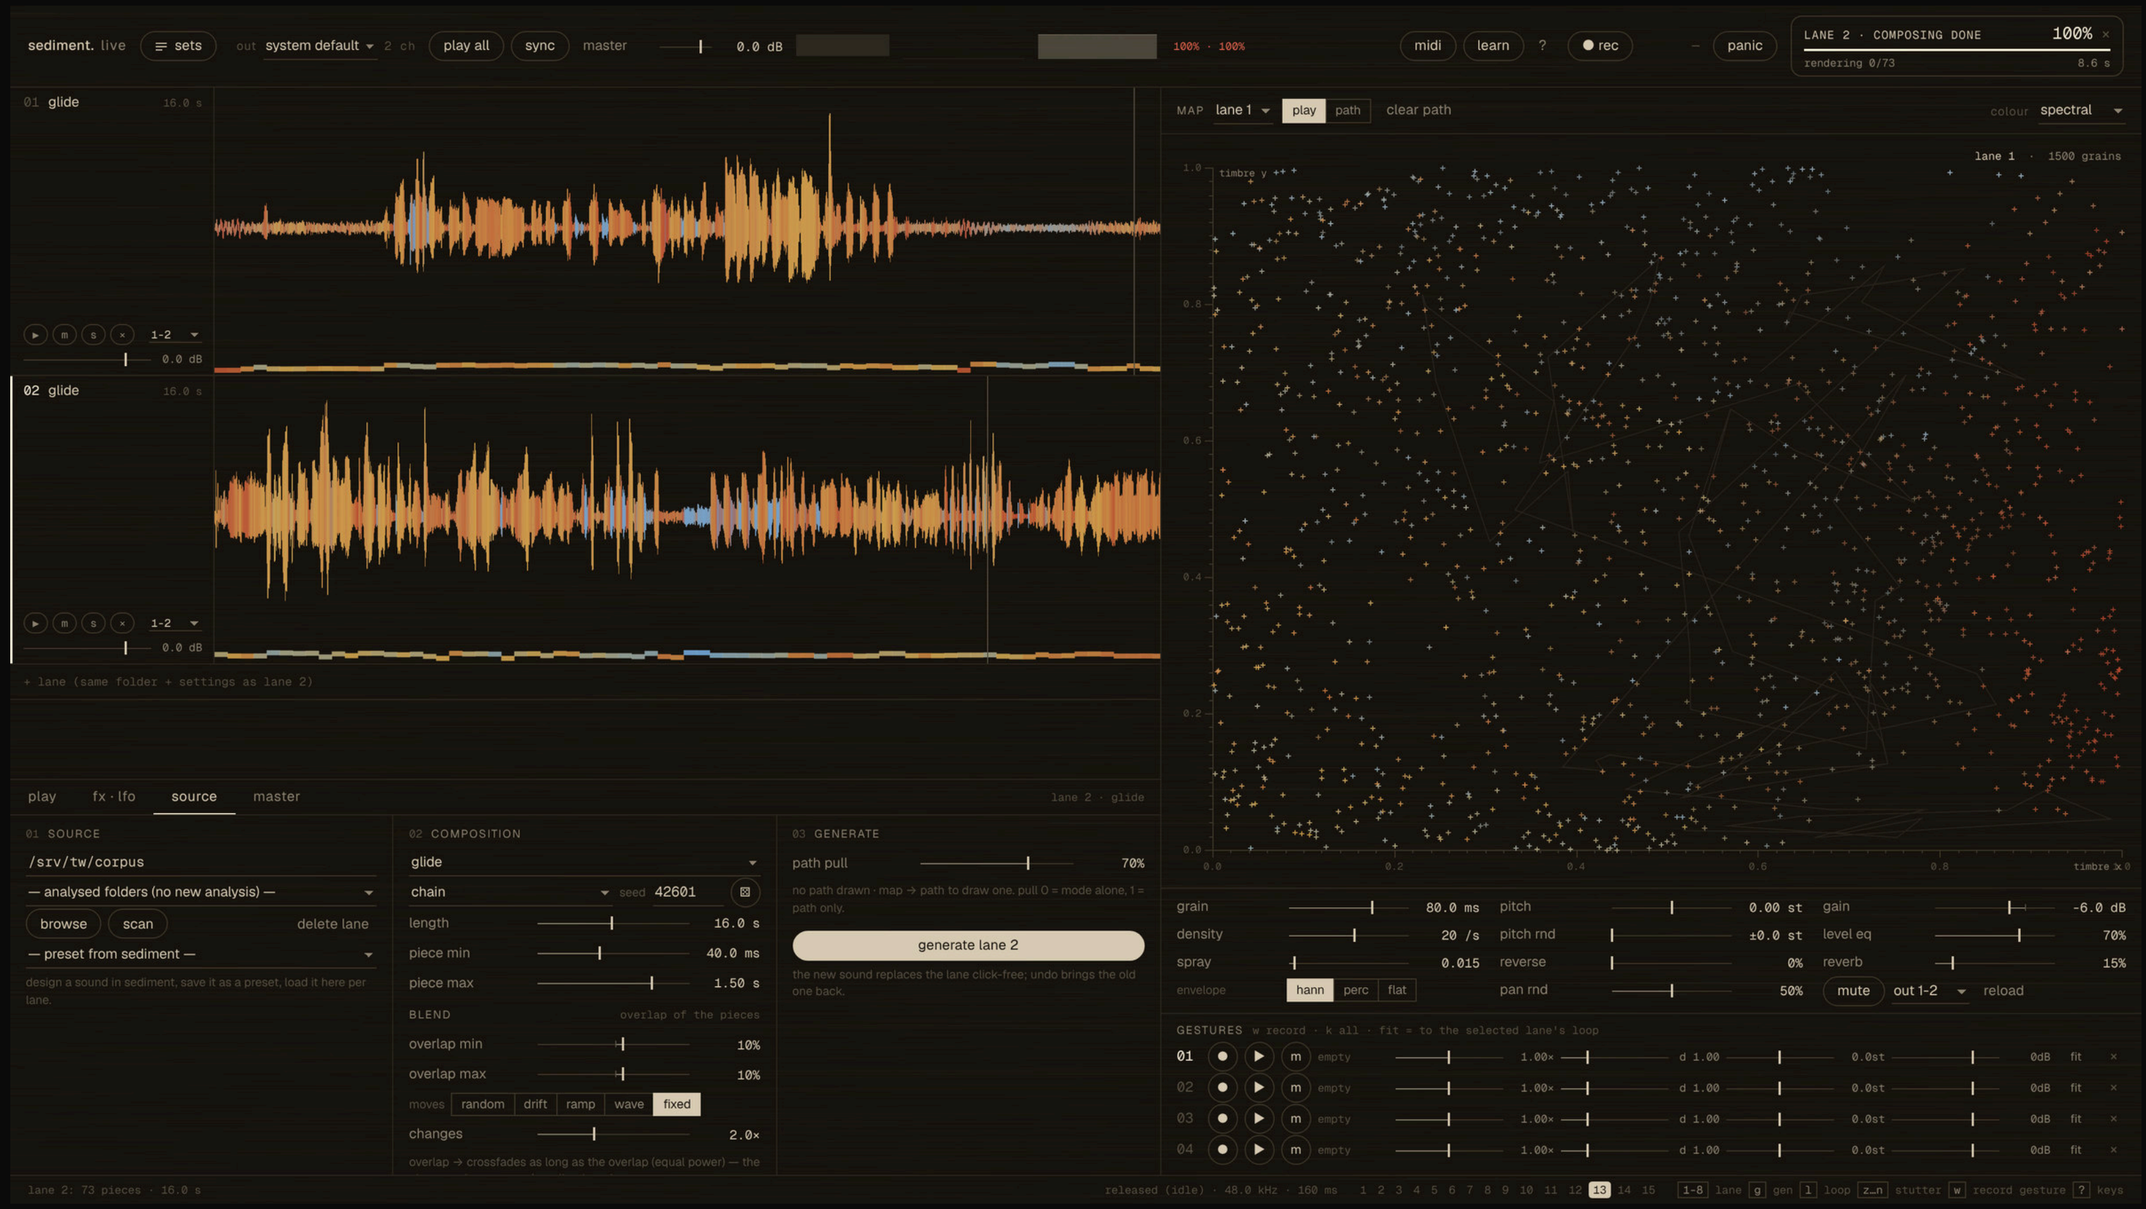
T02

sediment · live

2026

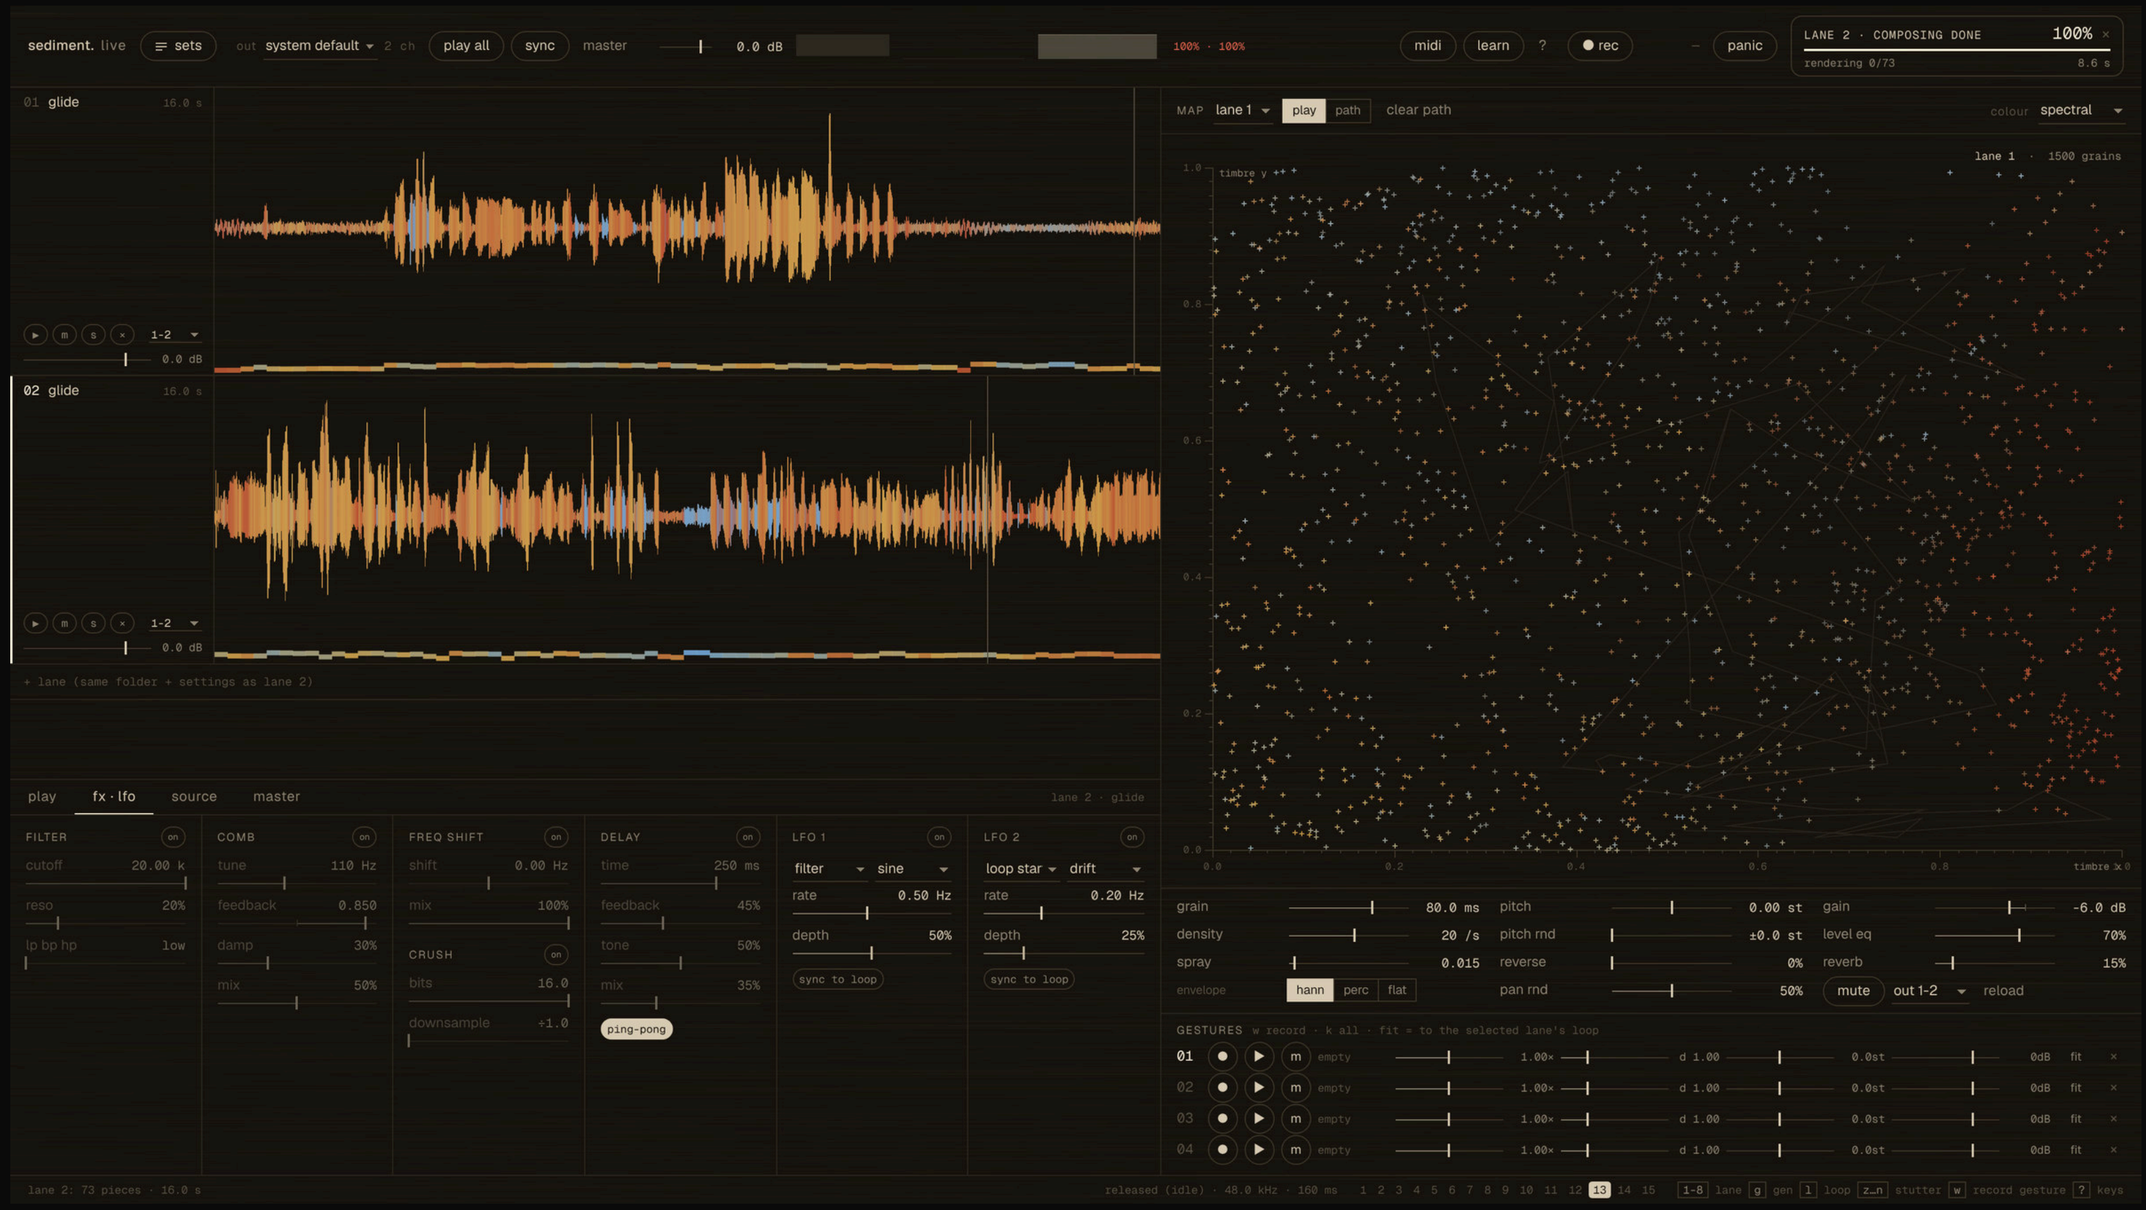
FIELD	Sound — corpus-based composition and live processing
VERSION	1.0
TECHNOLOGY	Python (aiohttp, sounddevice), Web Audio/browser UI, MIDI
SOURCE	sediment_live.zip

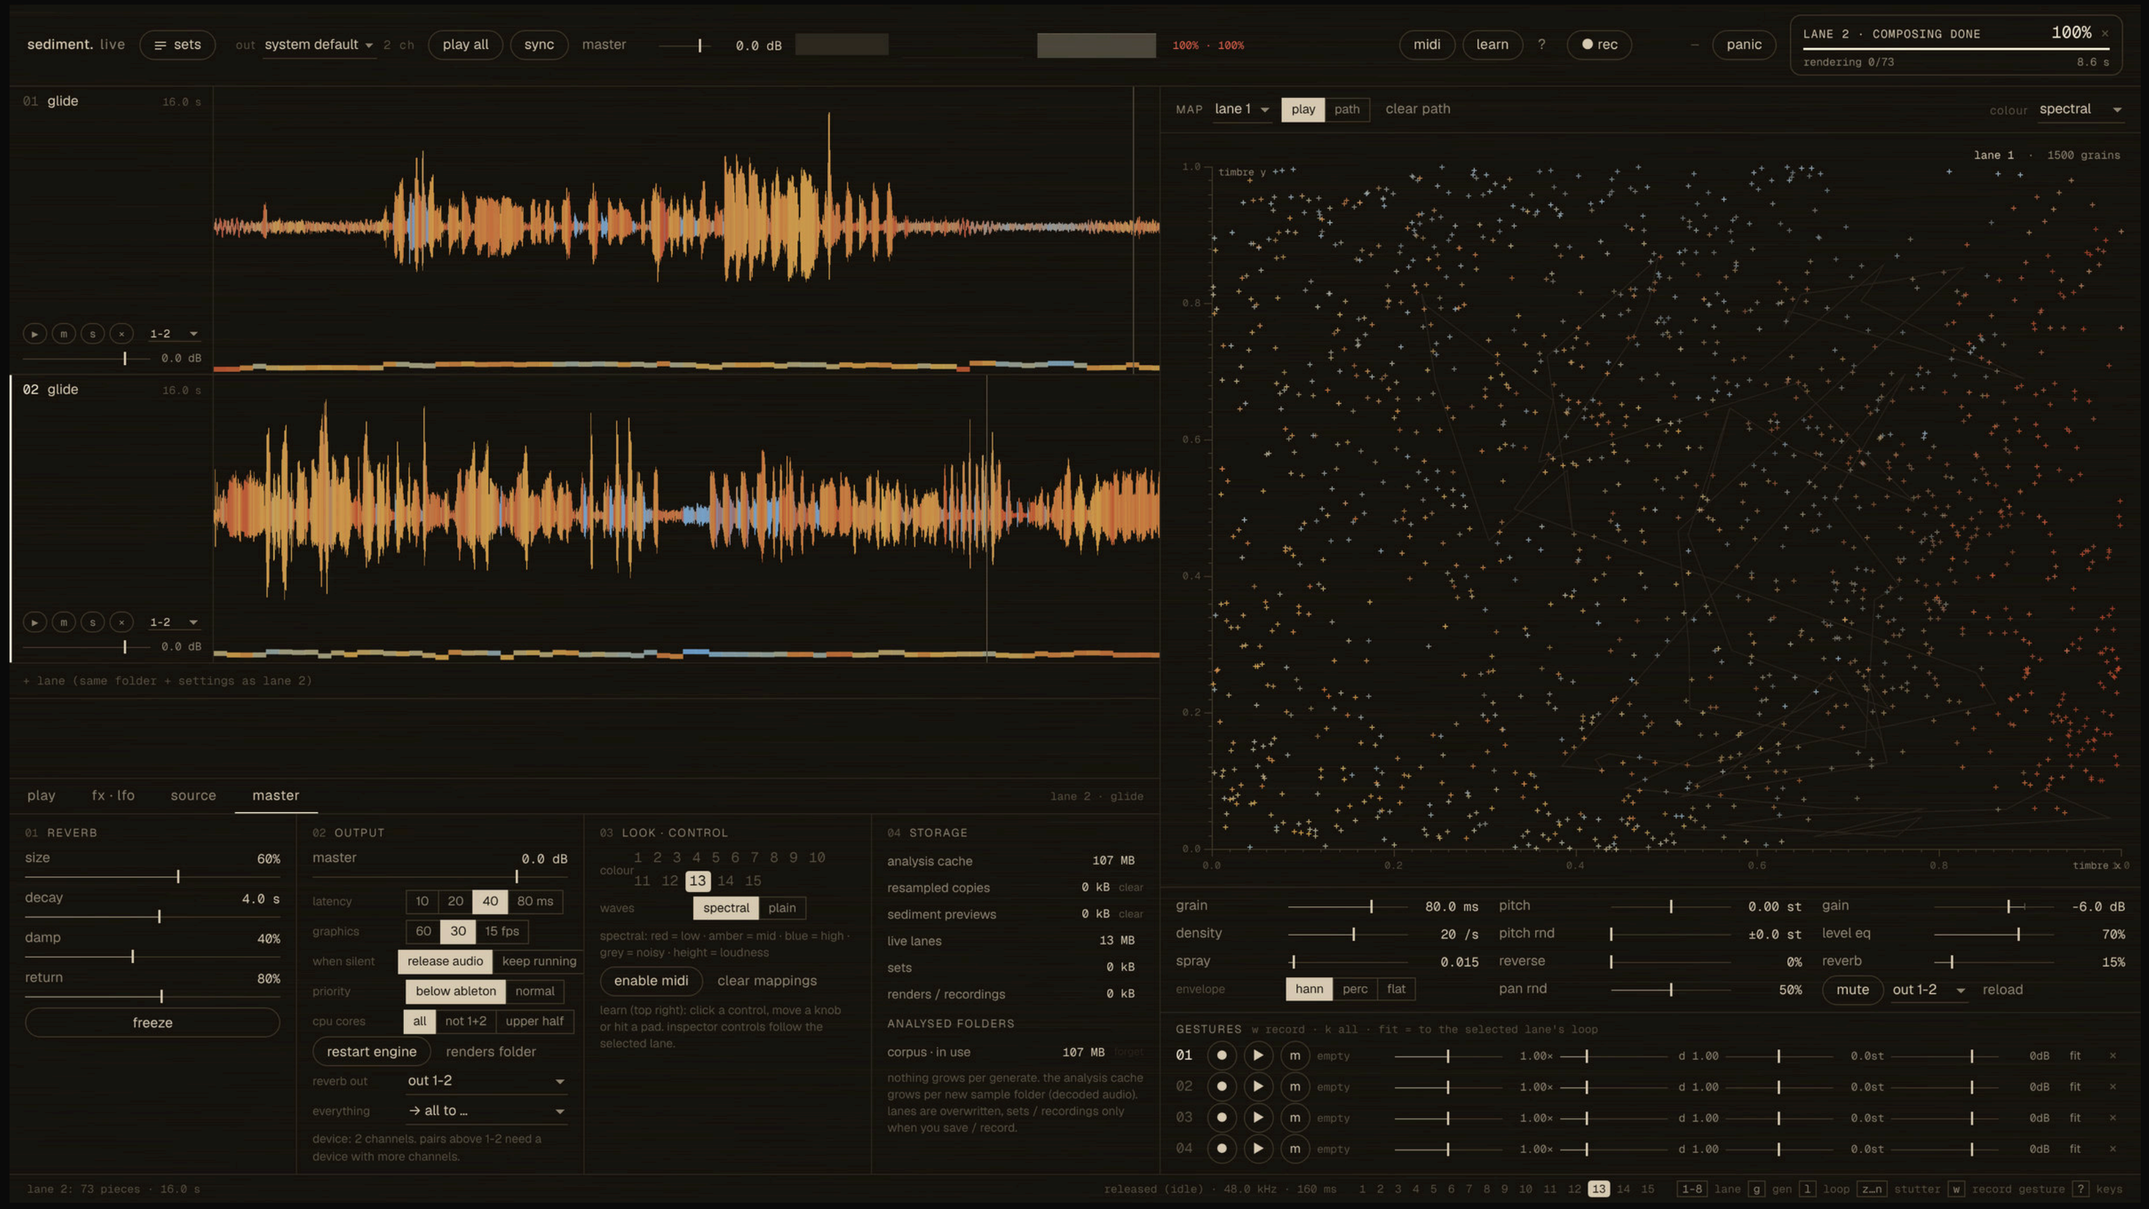
The performance branch of sediment. Up to eight lanes hold composed material that can be looped, stuttered, re-sorted and transformed in real time; the timbre map of each lane becomes a playable surface on which gestures are drawn, recorded and recalled. Sets store audio and state, MIDI learn binds every control. Designed to run alongside a DAW in concert.

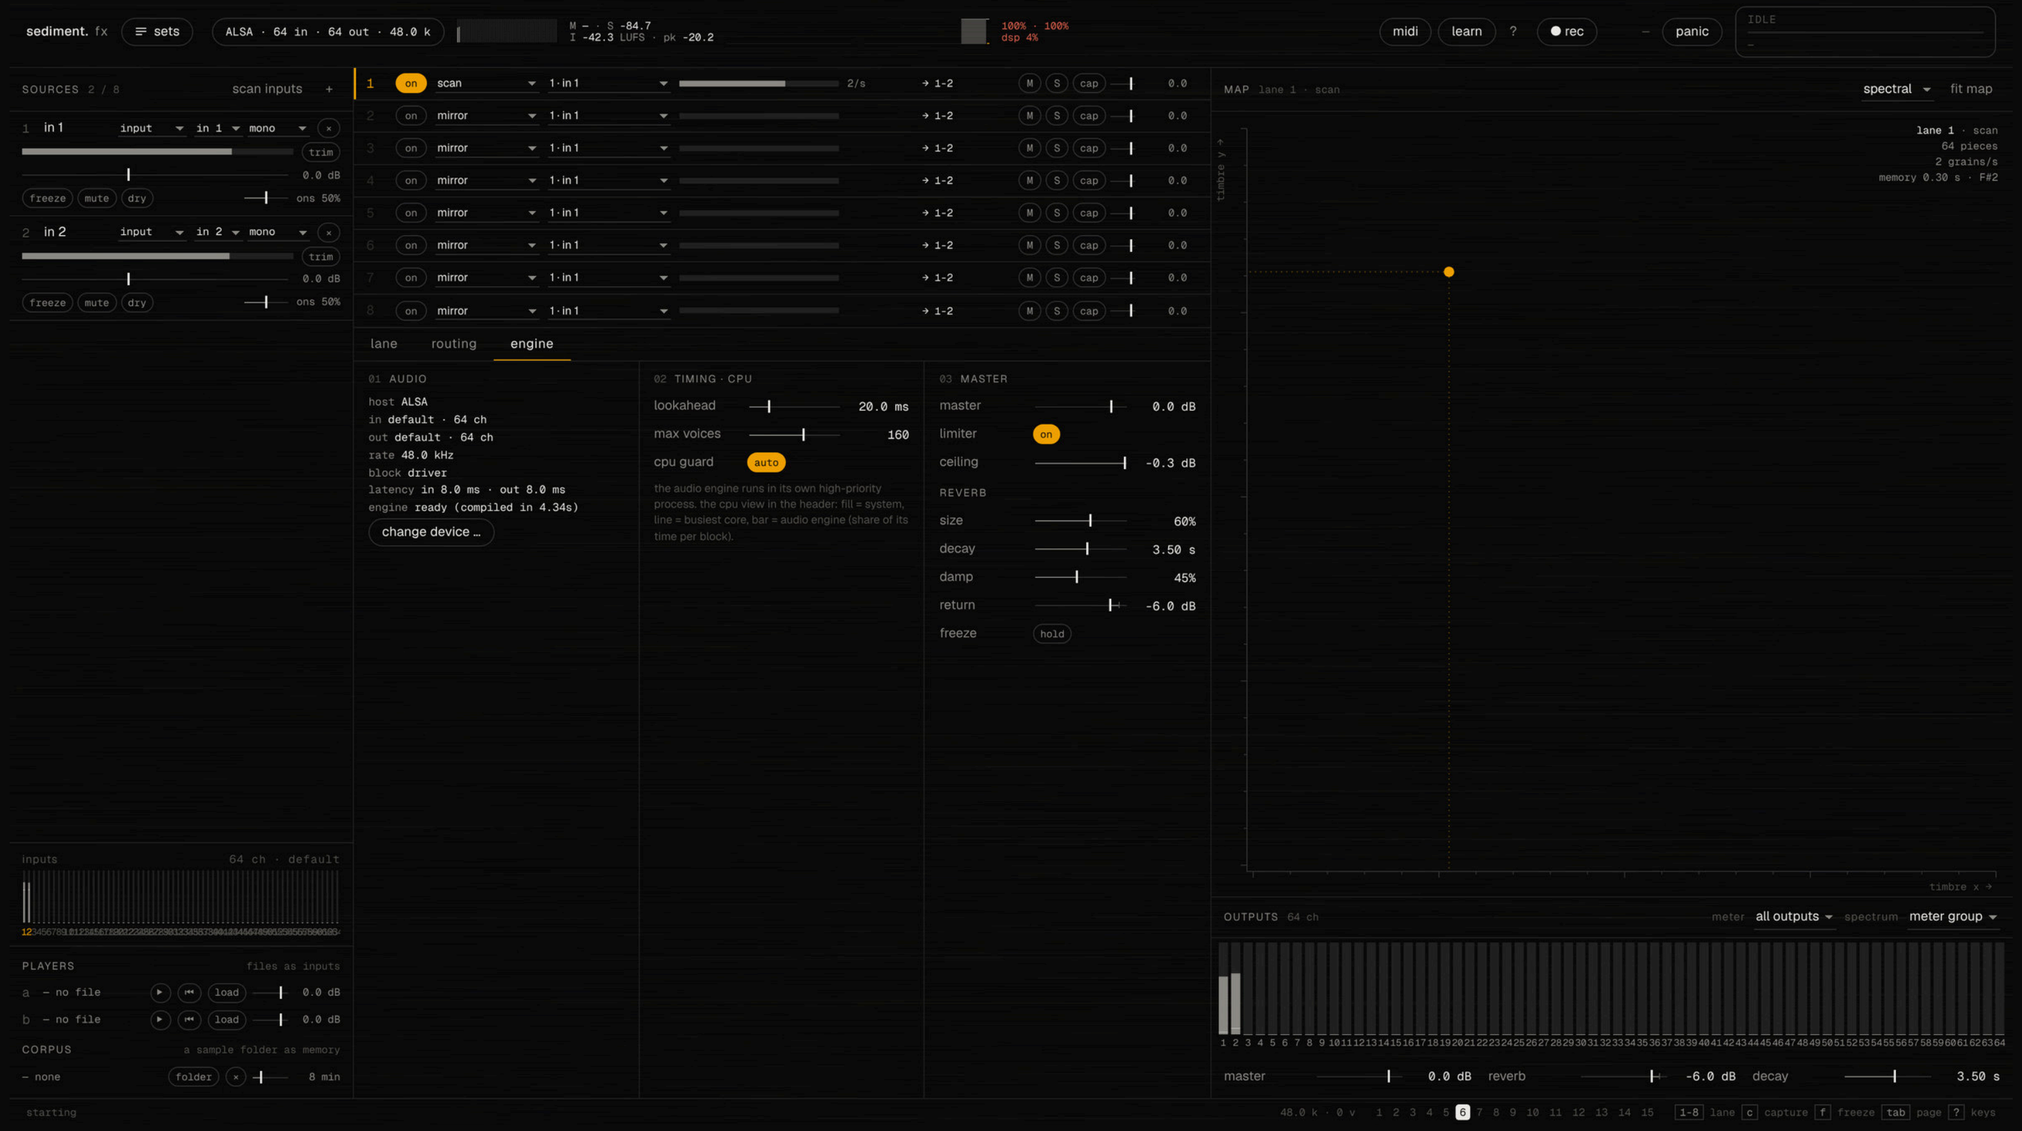


02 SOURCE PAGE

T02 · SEDIMENT · LIVE


sediment. live
sets
out system default 2 ch
play all
sync
master
0.0 dB
100% · 100%
midi
learn
?
rec
panic
LANE 2 · COMPOSING DONE 100%
rendering 0/73 8.6 s
01 glide 16.0 s
1-2
0.0 dB
02 glide 16.0 s
1-2
0.0 dB
+ lane (same folder + settings as lane 2)
MAP lane 1
play
path
clear path
colour spectral
lane 1 · 1500 grains
timbre y
timbre x
play
fx · lfo
source
master
lane 2 · glide
FILTER
cutoff 20.00 k
reso 20%
lp bp hp low
COMB
tune 110 Hz
feedback 0.850
damp 30%
mix 50%
FREQ SHIFT
shift 0.00 Hz
mix 100%
CRUSH
bits 16.0
downsample ÷1.0
DELAY
time 250 ms
feedback 45%
tone 50%
mix 35%
ping-pong
LFO 1
filter sine
rate 0.50 Hz
depth 50%
sync to loop
LFO 2
loop star drift
rate 0.20 Hz
depth 25%
sync to loop
grain 80.0 ms
density 20 /s
spray 0.015
envelope hann perc flat
pitch 0.00 st
pitch rnd ±0.0 st
reverse 0%
pan rnd 50%
gain -6.0 dB
level eq 70%
reverb 15%
mute
out 1-2
reload
GESTURES w record · k all · fit = to the selected lane's loop
01 empty 1.00× d 1.00 0.0st 0dB fit
02 empty 1.00× d 1.00 0.0st 0dB fit
03 empty 1.00× d 1.00 0.0st 0dB fit
04 empty 1.00× d 1.00 0.0st 0dB fit
lane 2: 73 pieces · 16.0 s
released (idle) · 48.0 kHz · 160 ms
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
1-8 lane g gen l loop z_n stutter w record gesture ? keys


sediment. live
sets
out system default 2 ch
play all
sync
master
0.0 dB
100% · 100%
midi
learn
?
rec
panic
LANE 2 · COMPOSING DONE 100%
rendering 0/73 8.6 s
01 glide 16.0 s
1-2
0.0 dB
02 glide 16.0 s
1-2
0.0 dB
+ lane (same folder + settings as lane 2)
MAP lane 1
play
path
clear path
colour spectral
lane 1 · 1500 grains
timbre y
timbre x
play
fx · lfo
source
master
lane 2 · glide
01 REVERB
size 60%
decay 4.0 s
damp 40%
return 80%
freeze
02 OUTPUT
master 0.0 dB
latency 10 20 40 80 ms
graphics 60 30 15 fps
when silent release audio keep running
priority below ableton normal
cpu cores all not 1+2 upper half
restart engine
renders folder
reverb out out 1-2
everything → all to …
device: 2 channels. pairs above 1-2 need a device with more channels.
03 LOOK · CONTROL
colour 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
waves spectral plain
spectral: red = low · amber = mid · blue = high · grey = noisy · height = loudness
enable midi
clear mappings
learn (top right): click a control, move a knob or hit a pad. inspector controls follow the selected lane.
04 STORAGE
analysis cache 107 MB
resampled copies 0 kB clear
sediment previews 0 kB clear
live lanes 13 MB
sets 0 kB
renders / recordings 0 kB
ANALYSED FOLDERS
corpus · in use 107 MB forget
nothing grows per generate. the analysis cache grows per new sample folder (decoded audio). lanes are overwritten, sets / recordings only when you save / record.
grain 80.0 ms
density 20 /s
spray 0.015
envelope hann perc flat
pitch 0.00 st
pitch rnd ±0.0 st
reverse 0%
pan rnd 50%
gain -6.0 dB
level eq 70%
reverb 15%
mute
out 1-2
reload
GESTURES w record · k all · fit = to the selected lane's loop
01 empty 1.00× d 1.00 0.0st 0dB fit
02 empty 1.00× d 1.00 0.0st 0dB fit
03 empty 1.00× d 1.00 0.0st 0dB fit
04 empty 1.00× d 1.00 0.0st 0dB fit
lane 2: 73 pieces · 16.0 s
released (idle) · 48.0 kHz · 160 ms
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
1–8 lane g gen l loop z…n stutter w record gesture ? keys


sediment. fx
sets
ALSA · 64 in · 64 out · 48.0 k
M – · S -84.7
I -42.3 LUFS · pk -20.2
100% · 100%
dsp 4%
midi
learn
?
rec
panic
IDLE
SOURCES 2 / 8
scan inputs
1 in 1 input in 1 mono trim 0.0 dB freeze mute dry ons 50%
2 in 2 input in 2 mono trim 0.0 dB freeze mute dry ons 50%
1 on scan 1·in 1 2/s → 1-2 M S cap 0.0
2 on mirror 1·in 1 → 1-2 M S cap 0.0
3 on mirror 1·in 1 → 1-2 M S cap 0.0
4 on mirror 1·in 1 → 1-2 M S cap 0.0
5 on mirror 1·in 1 → 1-2 M S cap 0.0
6 on mirror 1·in 1 → 1-2 M S cap 0.0
7 on mirror 1·in 1 → 1-2 M S cap 0.0
8 on mirror 1·in 1 → 1-2 M S cap 0.0
lane routing engine
01 AUDIO
host ALSA
in default · 64 ch
out default · 64 ch
rate 48.0 kHz
block driver
latency in 8.0 ms · out 8.0 ms
engine ready (compiled in 4.34s)
change device …
02 TIMING · CPU
lookahead 20.0 ms
max voices 160
cpu guard auto
the audio engine runs in its own high-priority process. the cpu view in the header: fill = system, line = busiest core, bar = audio engine (share of its time per block).
03 MASTER
master 0.0 dB
limiter on
ceiling -0.3 dB
REVERB
size 60%
decay 3.50 s
damp 45%
return -6.0 dB
freeze hold
MAP lane 1 · scan
spectral
fit map
lane 1 · scan
64 pieces
2 grains/s
memory 0.30 s · F#2
timbre ↑ bright
timbre x →
inputs 64 ch · default
PLAYERS files as inputs
a – no file load 0.0 dB
b – no file load 0.0 dB
CORPUS a sample folder as memory
– none folder 8 min
OUTPUTS 64 ch
meter all outputs
spectrum meter group
master 0.0 dB
reverb -6.0 dB
decay 3.50 s
starting
48.0 k · 0 v
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
1-8 lane c capture f freeze tab page ? keys

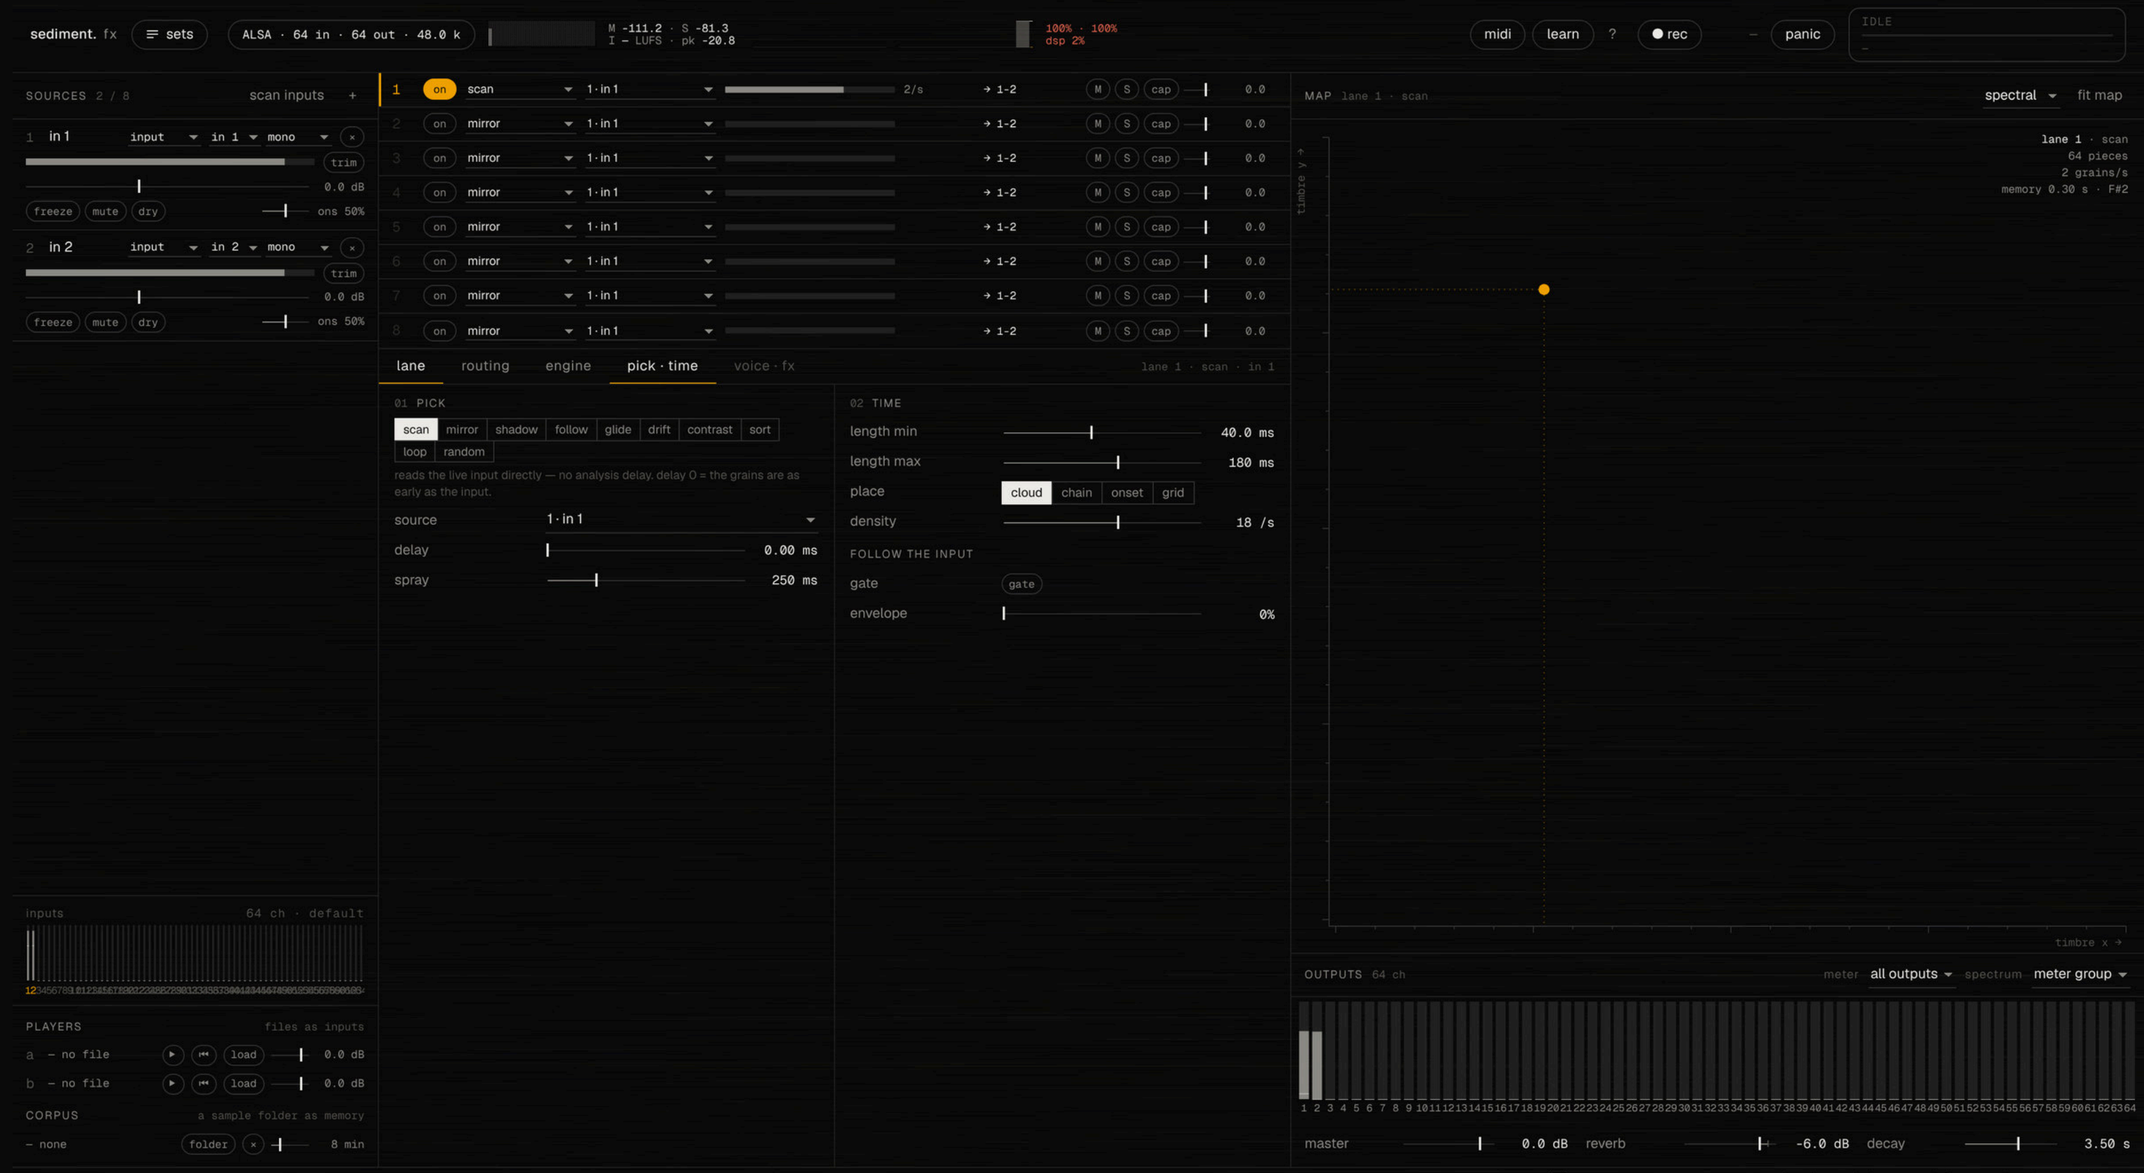
T03

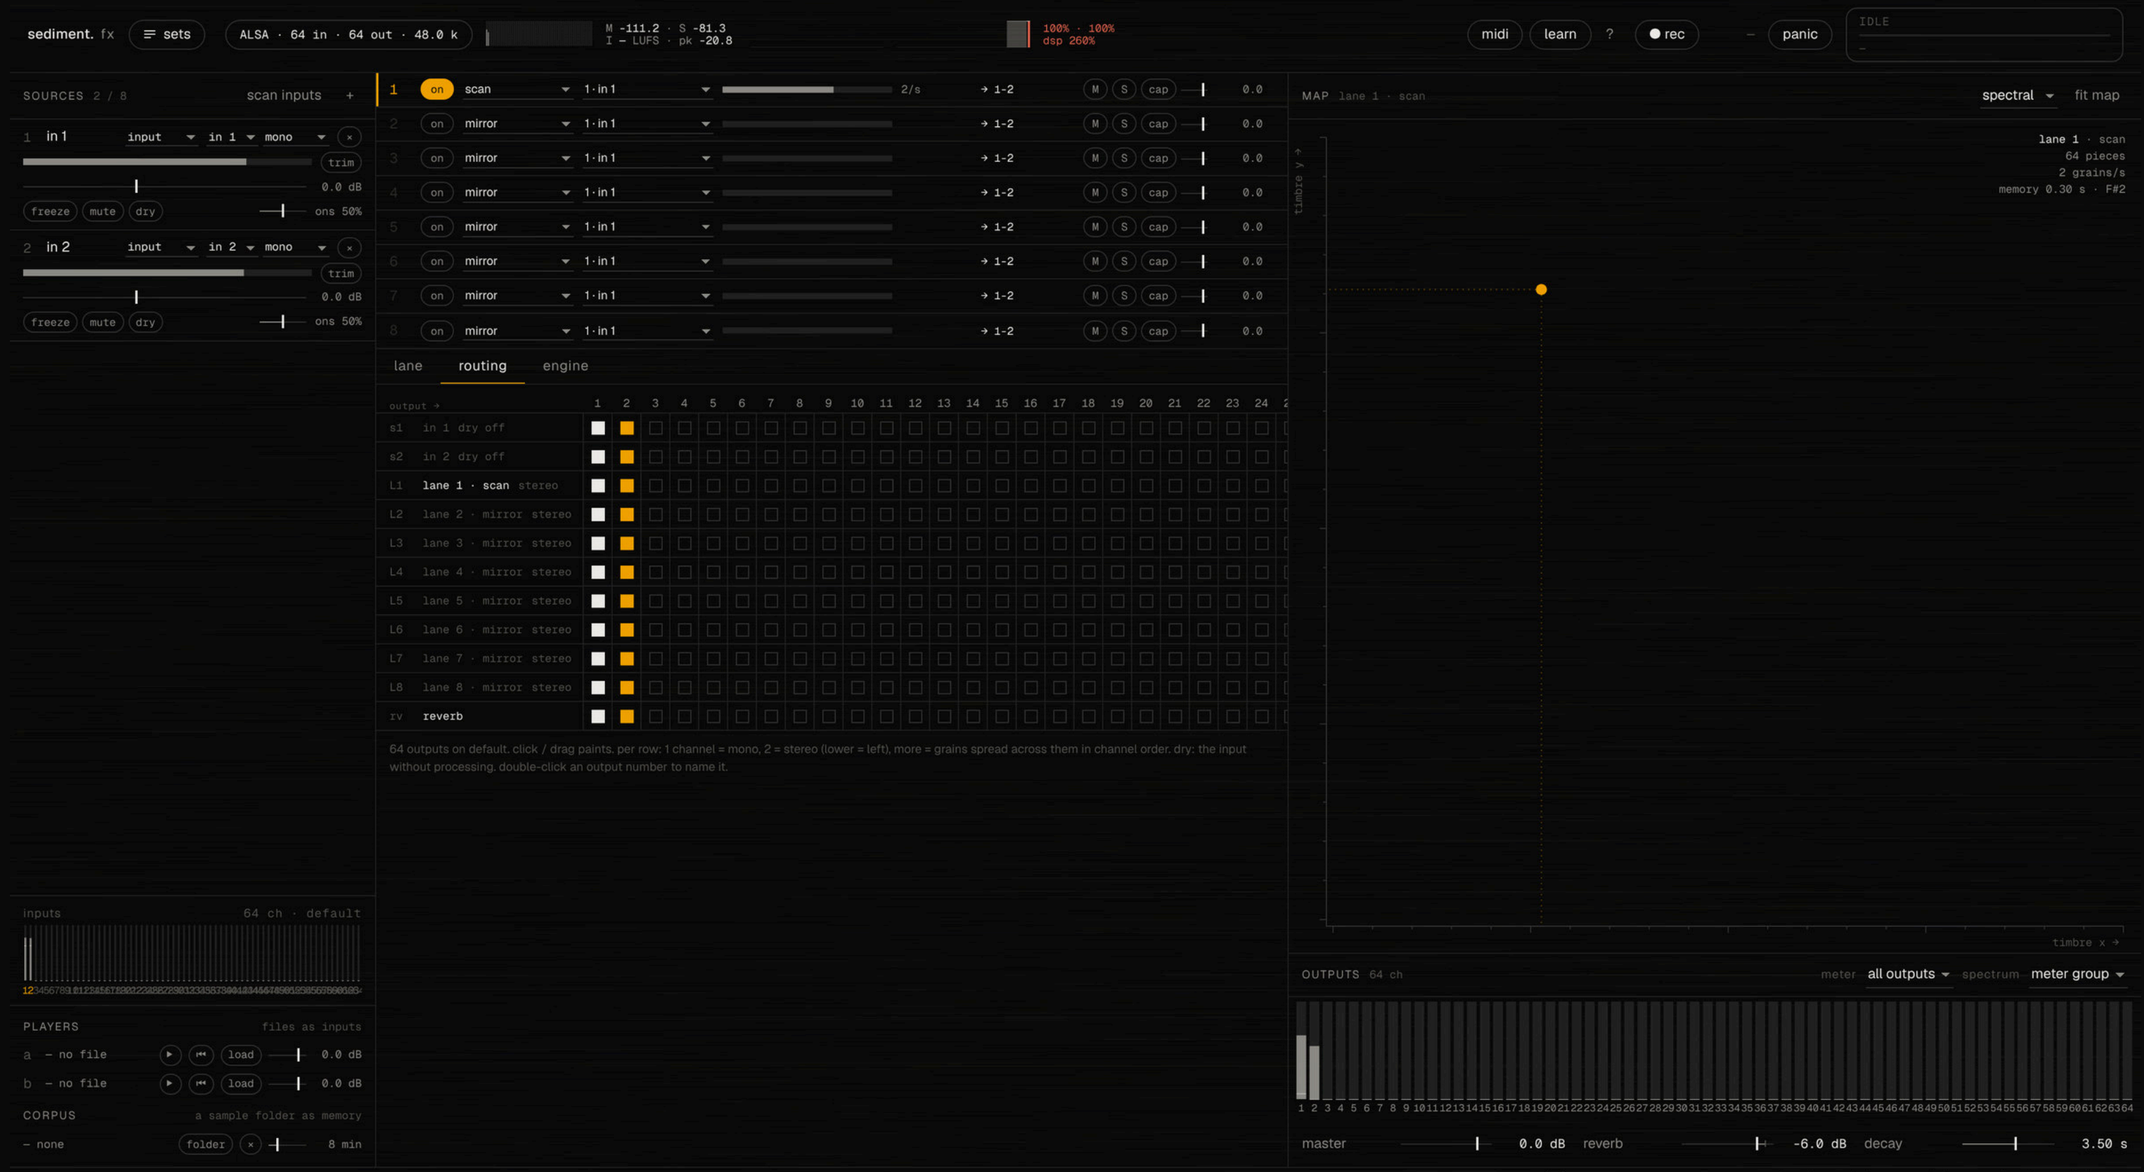
sediment · fx

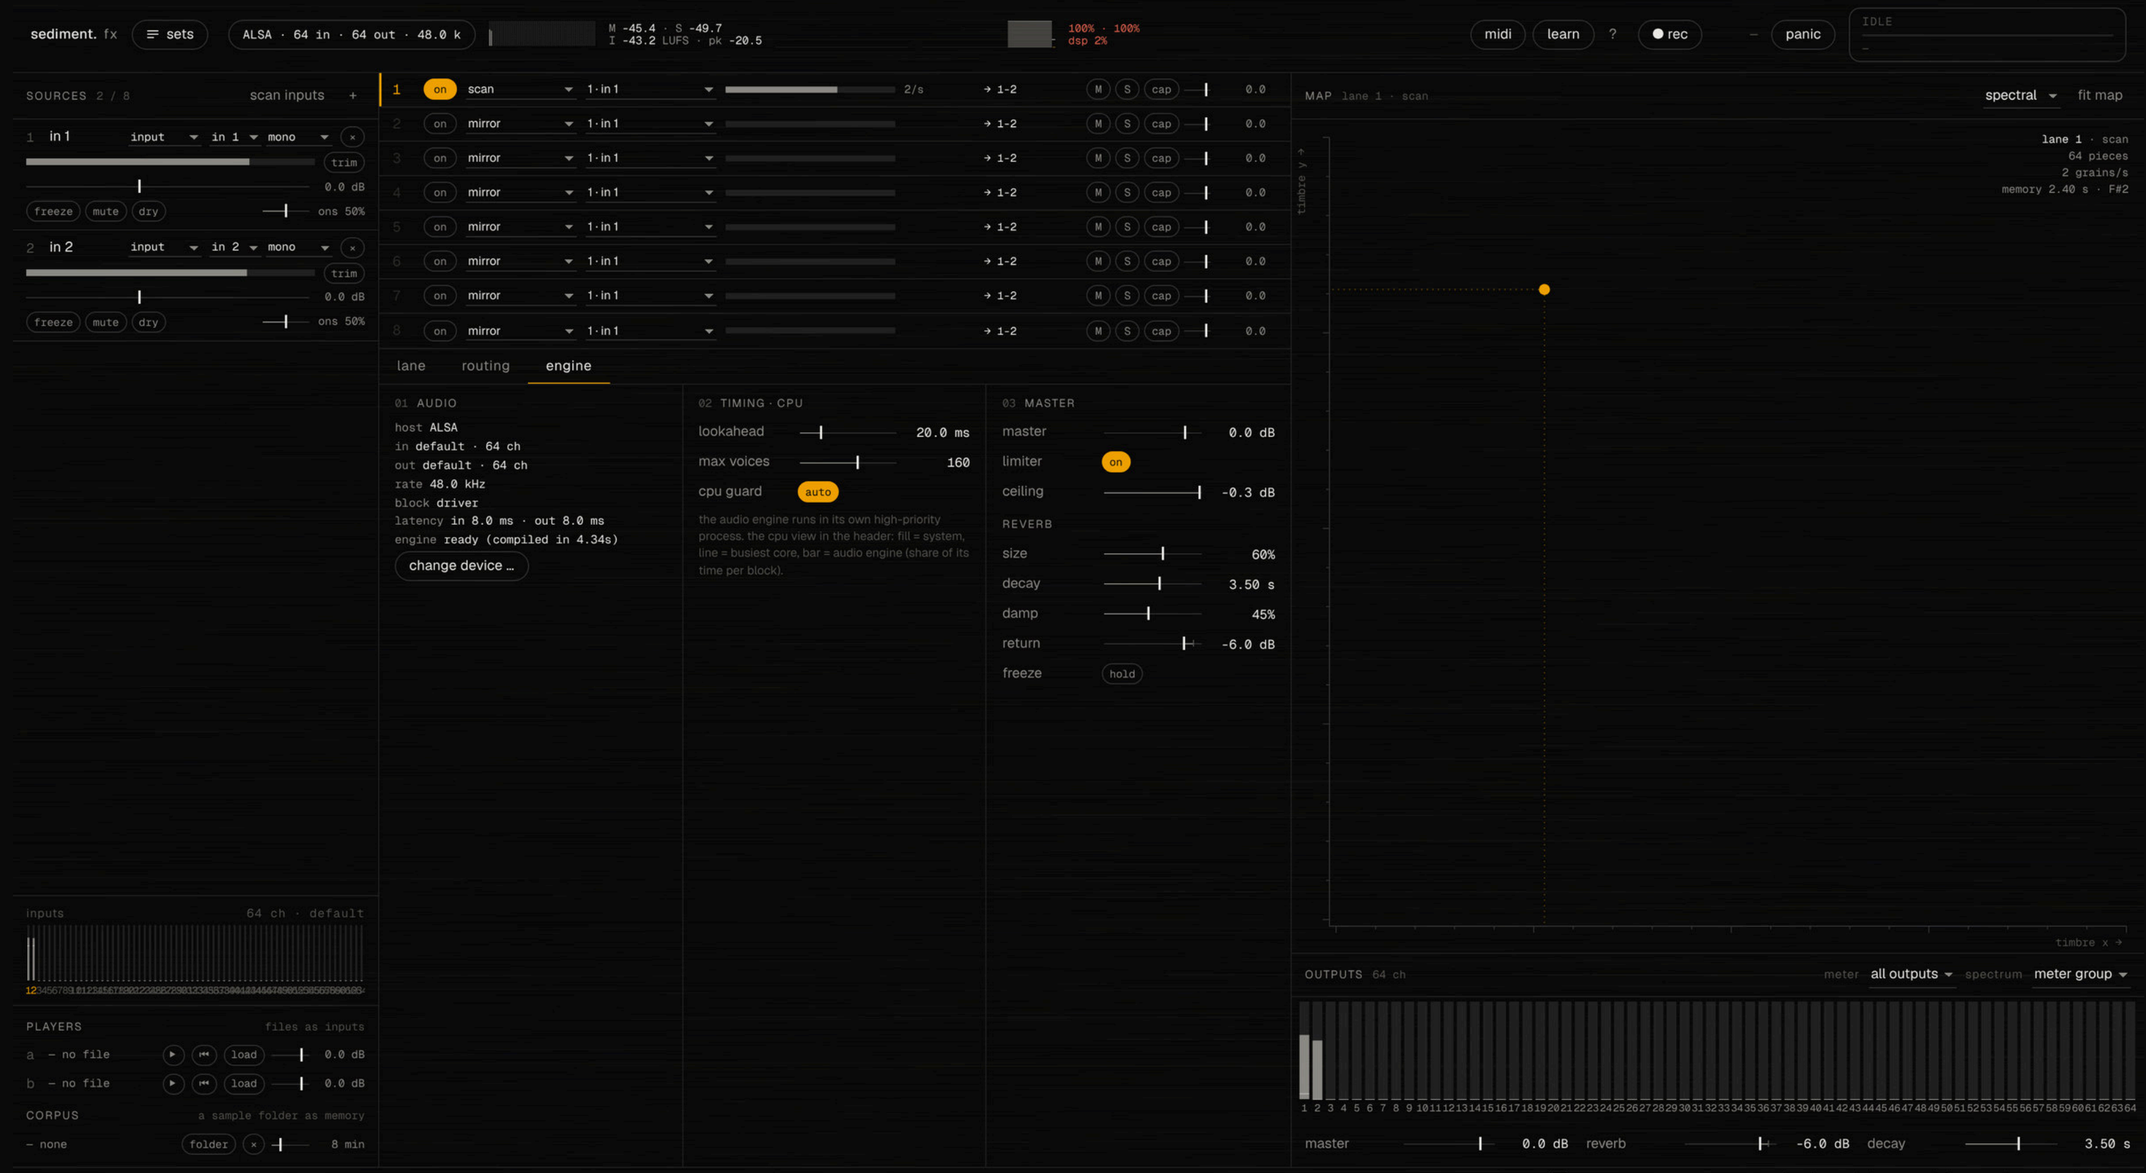
2026

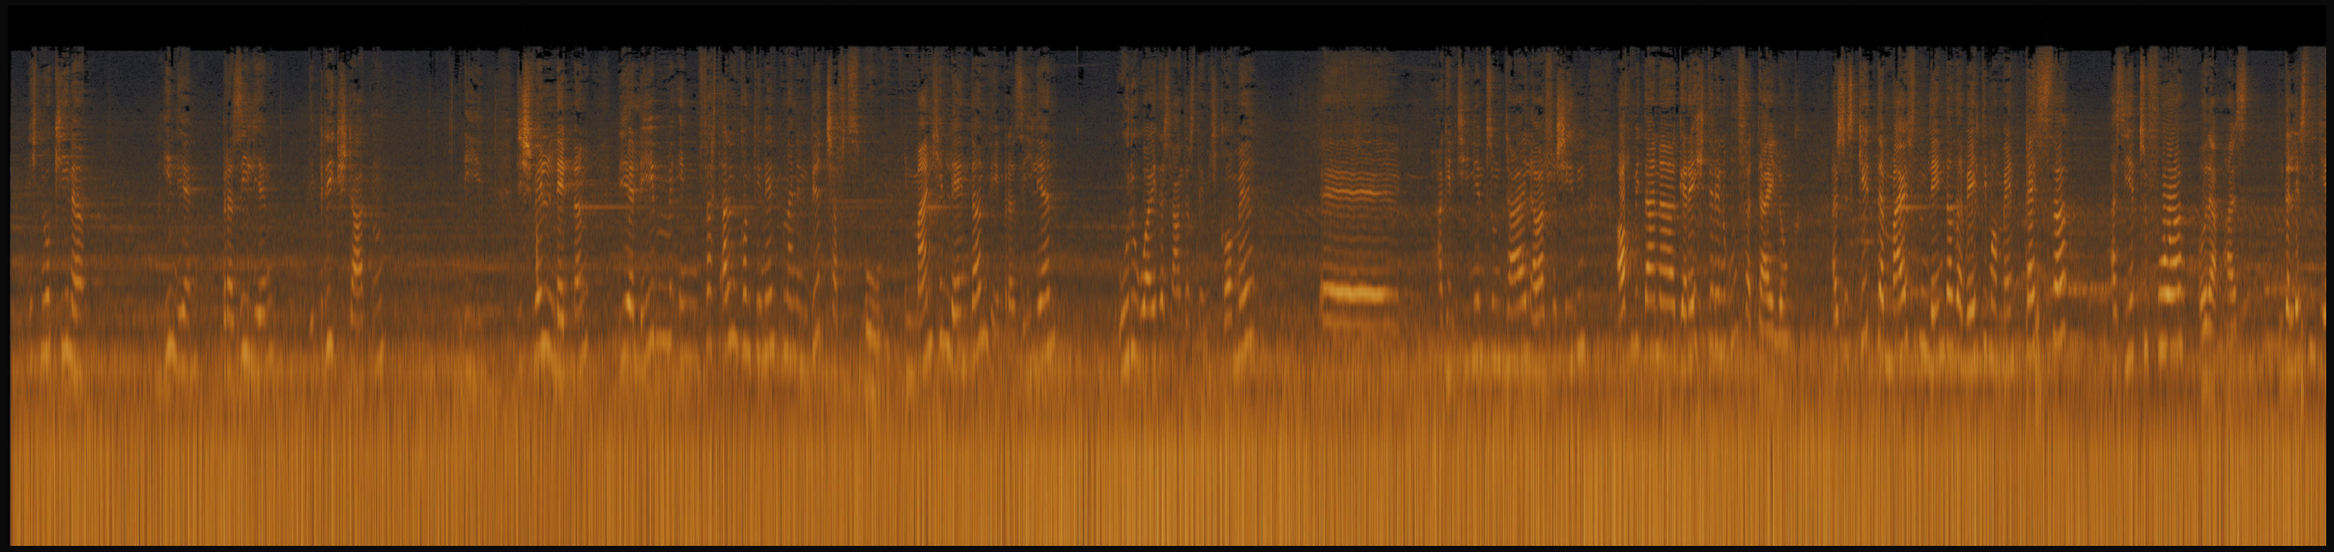
FIELD	Sound — corpus-based composition and live processing
VERSION	1.0
TECHNOLOGY	Python, numba-compiled audio engine, shared-memory audio process, ASIO/WASAPI/CoreAudio
SOURCE	sediment_fx_1.0.zip

A realtime effect unit applying sediment's ordering modes to live input. Incoming audio is written to memory and analysed as it arrives; up to eight lanes re-order pieces of the immediate past in time and route them to arbitrary discrete outputs. Architecturally it separates a compiled, allocation-free audio body from an analysing ‘brain’ that schedules events ahead, which makes multichannel operation at low latency possible.


sediment. fx
sets
ALSA · 64 in · 64 out · 48.0 k
M -111.2 · S -81.3
I — LUFS · pk -20.8
100% · 100%
dsp 2%
midi
learn
?
rec
panic
IDLE
SOURCES 2 / 8
scan inputs
1 in 1 input in 1 mono
trim
0.0 dB
freeze mute dry
ons 50%
2 in 2 input in 2 mono
trim
0.0 dB
freeze mute dry
ons 50%
1 on scan 1 · in 1 2/s → 1-2 M S cap 0.0
2 on mirror 1 · in 1 → 1-2 M S cap 0.0
3 on mirror 1 · in 1 → 1-2 M S cap 0.0
4 on mirror 1 · in 1 → 1-2 M S cap 0.0
5 on mirror 1 · in 1 → 1-2 M S cap 0.0
6 on mirror 1 · in 1 → 1-2 M S cap 0.0
7 on mirror 1 · in 1 → 1-2 M S cap 0.0
8 on mirror 1 · in 1 → 1-2 M S cap 0.0
lane routing engine pick · time voice · fx
lane 1 · scan · in 1
01 PICK
scan mirror shadow follow glide drift contrast sort
loop random
reads the live input directly — no analysis delay. delay 0 = the grains are as early as the input.
source 1 · in 1
delay 0.00 ms
spray 250 ms
02 TIME
length min 40.0 ms
length max 180 ms
place cloud chain onset grid
density 18 /s
FOLLOW THE INPUT
gate gate
envelope 0%
MAP lane 1 · scan
spectral
fit map
lane 1 · scan
64 pieces
2 grains/s
memory 0.30 s · F#2
timbre y →
timbre x →
inputs
64 ch · default
PLAYERS
files as inputs
a — no file load 0.0 dB
b — no file load 0.0 dB
CORPUS
a sample folder as memory
— none folder 8 min
OUTPUTS 64 ch
meter all outputs
spectrum meter group
master 0.0 dB
reverb -6.0 dB
decay 3.50 s
starting
48.0 k · 0 v
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
1-8 lane c capture f freeze tab page ? keys


sediment. fx
sets
ALSA · 64 in · 64 out · 48.0 k
M -111.2 · S -81.3
I – LUFS · pk -20.8
100% · 100%
dsp 260%
midi
learn
?
rec
panic
IDLE
SOURCES 2 / 8
scan inputs
1 in 1 input in 1 mono
trim
0.0 dB
freeze mute dry
ons 50%
2 in 2 input in 2 mono
trim
0.0 dB
freeze mute dry
ons 50%
1 on scan 1 · in 1 2/s → 1-2 M S cap 0.0
2 on mirror 1 · in 1 → 1-2 M S cap 0.0
3 on mirror 1 · in 1 → 1-2 M S cap 0.0
4 on mirror 1 · in 1 → 1-2 M S cap 0.0
5 on mirror 1 · in 1 → 1-2 M S cap 0.0
6 on mirror 1 · in 1 → 1-2 M S cap 0.0
7 on mirror 1 · in 1 → 1-2 M S cap 0.0
8 on mirror 1 · in 1 → 1-2 M S cap 0.0
lane routing engine
output →
s1 in 1 dry off
s2 in 2 dry off
L1 lane 1 · scan stereo
L2 lane 2 · mirror stereo
L3 lane 3 · mirror stereo
L4 lane 4 · mirror stereo
L5 lane 5 · mirror stereo
L6 lane 6 · mirror stereo
L7 lane 7 · mirror stereo
L8 lane 8 · mirror stereo
rv reverb
64 outputs on default. click / drag paints. per row: 1 channel = mono, 2 = stereo (lower = left), more = grains spread across them in channel order. dry: the input without processing. double-click an output number to name it.
MAP lane 1 · scan
spectral
fit map
lane 1 · scan
64 pieces
2 grains/s
memory 0.30 s · F#2
timbre y →
timbre x →
inputs
64 ch · default
PLAYERS
files as inputs
a – no file load 0.0 dB
b – no file load 0.0 dB
CORPUS
a sample folder as memory
– none folder 8 min
OUTPUTS 64 ch
meter all outputs
spectrum meter group
master 0.0 dB
reverb -6.0 dB
decay 3.50 s
starting
48.0 k · 1 v
1-8 lane
c capture
f freeze
tab page
? keys


sediment. fx
sets
ALSA · 64 in · 64 out · 48.0 k
M -45.4 · S -49.7
I -43.2 LUFS · pk -20.5
100% · 100%
dsp 2%
midi
learn
?
rec
panic
IDLE
SOURCES 2 / 8
scan inputs
1 in 1 input in 1 mono
trim
0.0 dB
freeze mute dry
ons 50%
2 in 2 input in 2 mono
trim
0.0 dB
freeze mute dry
ons 50%
1 on scan 1 · in 1 2/s → 1-2 M S cap 0.0
2 on mirror 1 · in 1 → 1-2 M S cap 0.0
3 on mirror 1 · in 1 → 1-2 M S cap 0.0
4 on mirror 1 · in 1 → 1-2 M S cap 0.0
5 on mirror 1 · in 1 → 1-2 M S cap 0.0
6 on mirror 1 · in 1 → 1-2 M S cap 0.0
7 on mirror 1 · in 1 → 1-2 M S cap 0.0
8 on mirror 1 · in 1 → 1-2 M S cap 0.0
lane
routing
engine
01 AUDIO
host ALSA
in default · 64 ch
out default · 64 ch
rate 48.0 kHz
block driver
latency in 8.0 ms · out 8.0 ms
engine ready (compiled in 4.34s)
change device …
02 TIMING · CPU
lookahead 20.0 ms
max voices 160
cpu guard auto
the audio engine runs in its own high-priority process. the cpu view in the header: fill = system, line = busiest core, bar = audio engine (share of its time per block).
03 MASTER
master 0.0 dB
limiter on
ceiling -0.3 dB
REVERB
size 60%
decay 3.50 s
damp 45%
return -6.0 dB
freeze hold
MAP lane 1 · scan
spectral
fit map
lane 1 · scan
64 pieces
2 grains/s
memory 2.40 s · F#2
timbre y →
timbre x →
inputs
64 ch · default
PLAYERS
files as inputs
a – no file load 0.0 dB
b – no file load 0.0 dB
CORPUS
a sample folder as memory
– none folder 8 min
OUTPUTS 64 ch
meter all outputs
spectrum meter group
master 0.0 dB
reverb -6.0 dB
decay 3.50 s
starting
48.0 k · 0 v
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
1-8 lane c capture f freeze tab page ? keys



sediment · pod

2026

SOUND

01 · input · radio recording, 30 s

T04

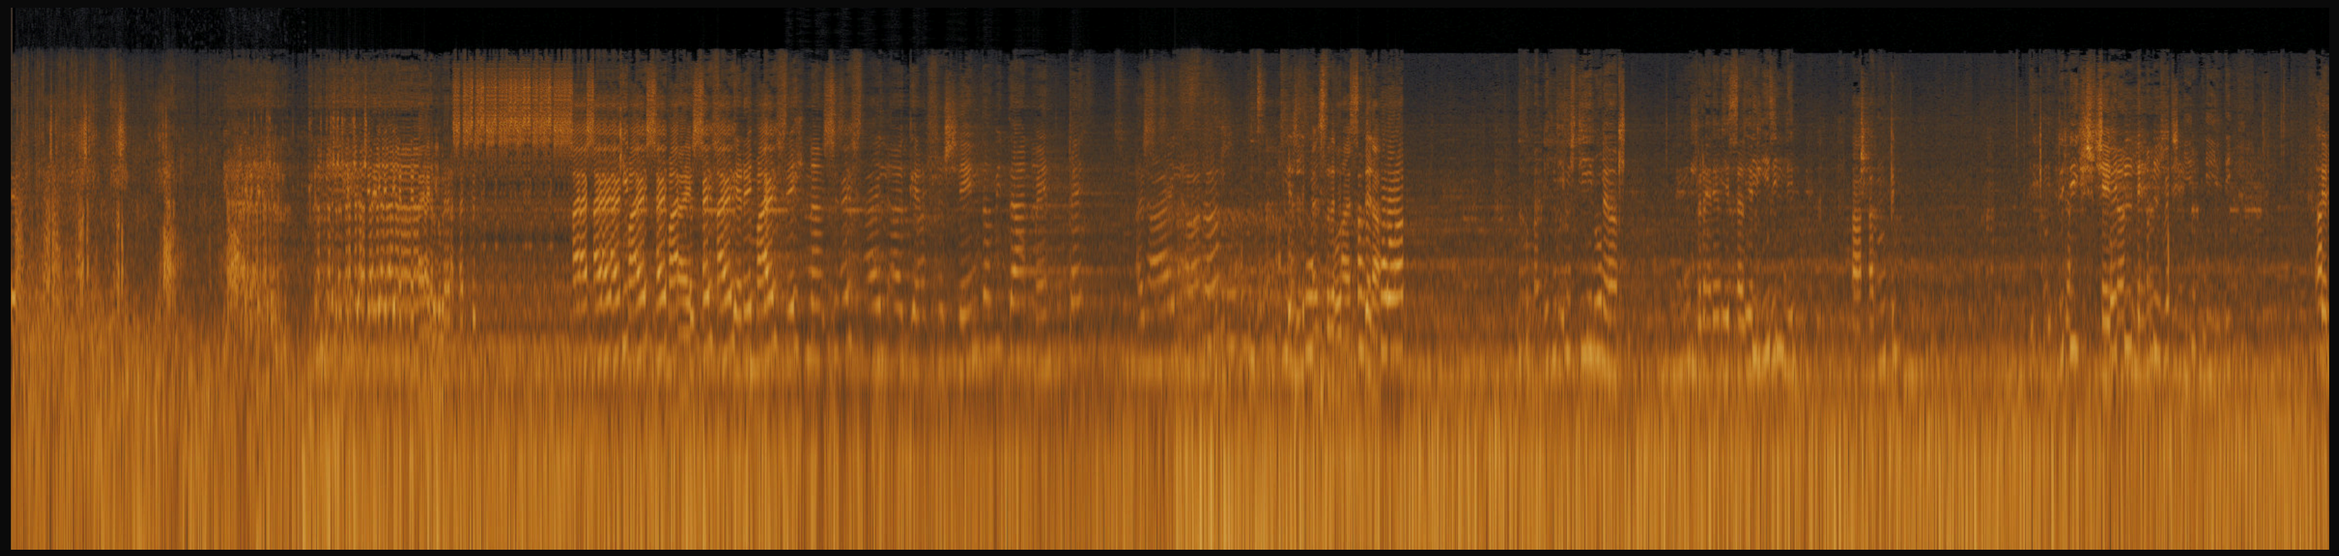
T04

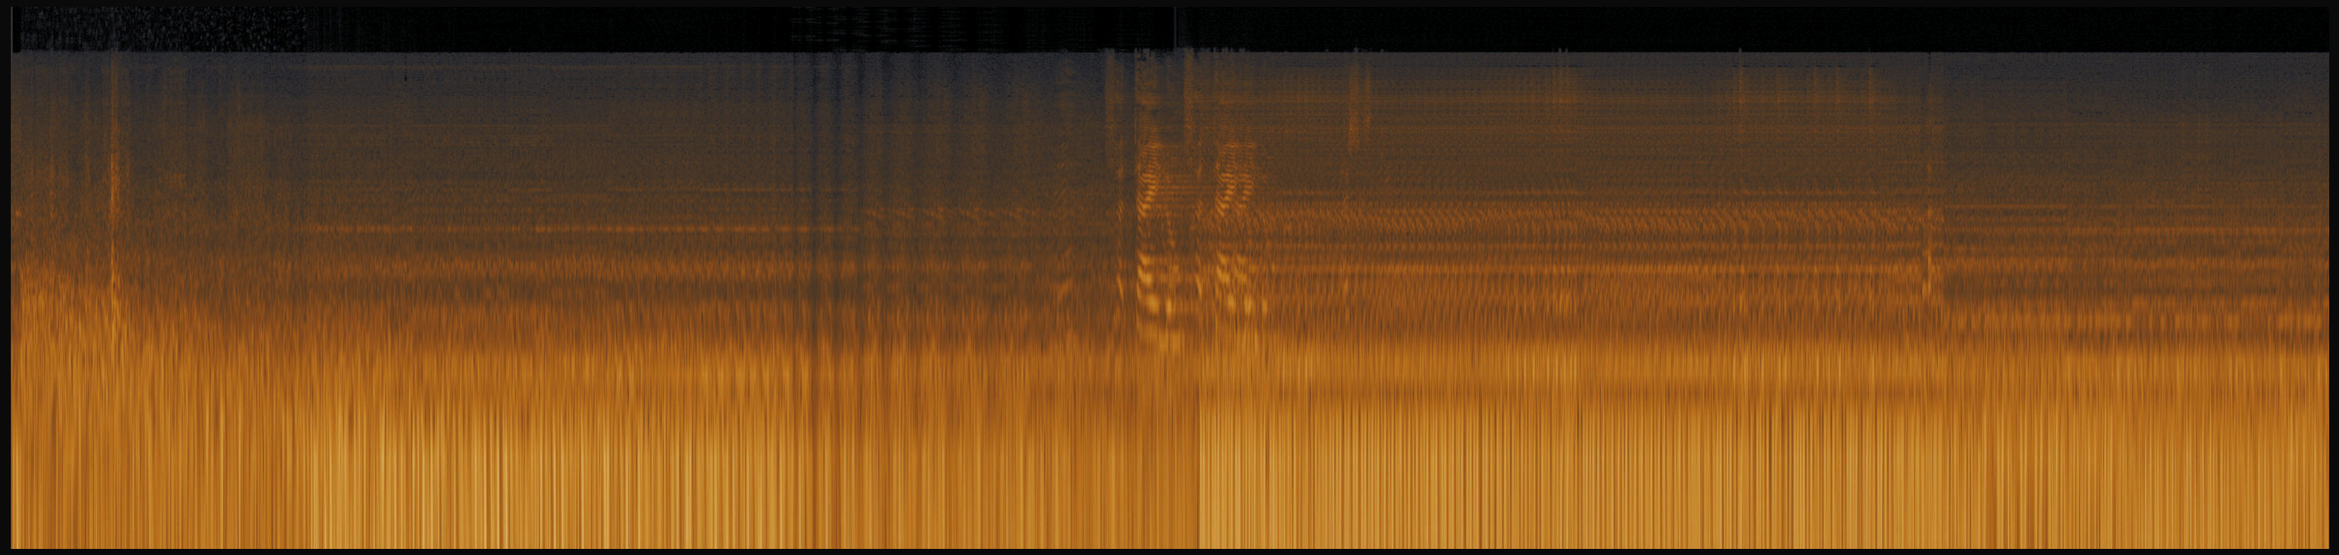
sediment · pod

2026

FIELD	Sound — corpus-based composition and live processing
VERSION	v3
TECHNOLOGY	C++ firmware for Electrosmith Daisy Pod (STM32H7); host-side offline renderer
SOURCE	sediment_pod_v3.zip

Firmware that ports sediment · fx to the Electrosmith Daisy Pod. Stereo line input is held in 87 s of memory; sixteen ‘moves’ (scan, loop, sorts by loudness, brightness, pitch, noise and flux, mirror, shadow, glide, contrast, drift, onset, grid) re-order pieces of it. One knob spans the grain continuum from 2 ms buzz to 2.5 s wash, the other controls the move. Plate: spectrograms rendered offline from the firmware's own engine on a radio recording.

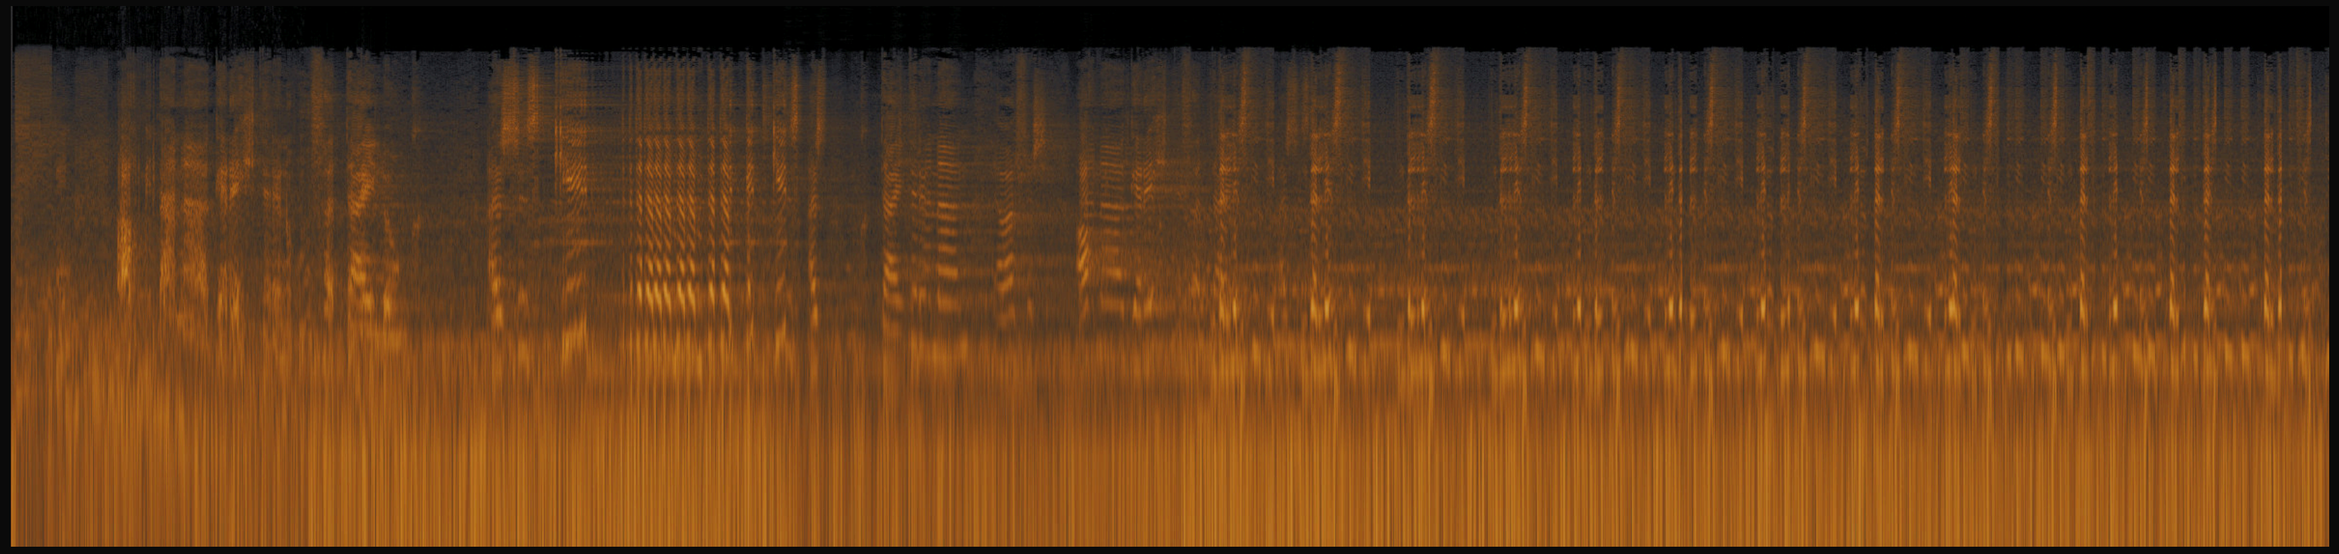


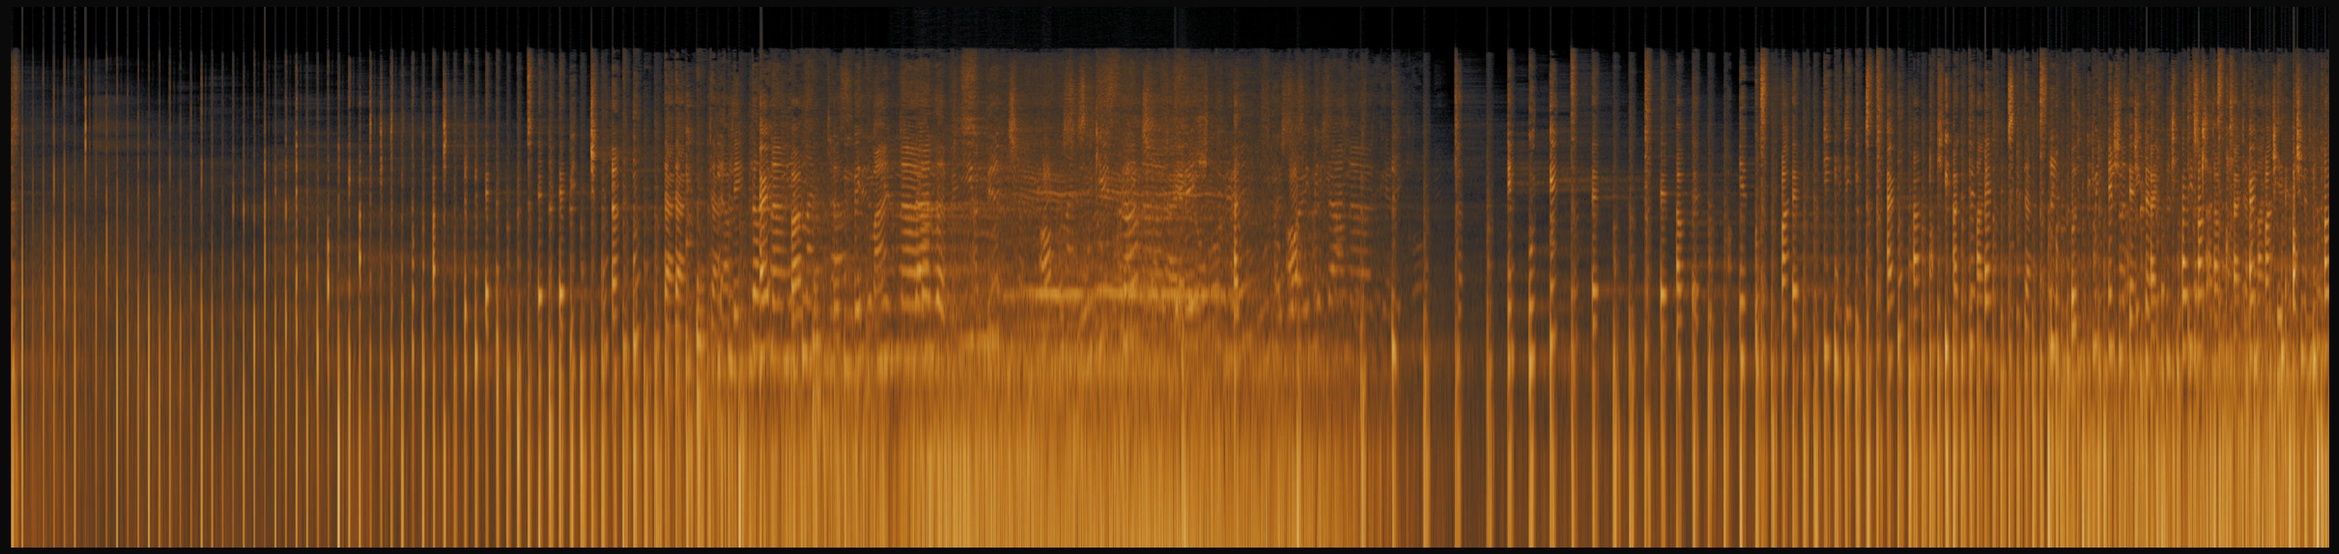


02 MOVE 5 · SORT BRIGHT

03 MOVE 11 · GLIDE

T04 · SEDIMENT · POD





04 MOVE 2 · LOOP 05 MOVE 16 · GRID

T04 · SEDIMENT · POD


ligature
video-bound scoring
▶ play
loop
00:07.20 / 00:32.00
undo
redo
final
render
cpu 100% engine 0% · 153MB
ui 0.0/0ms · 0 audio idle
render soundtrack · failed
load a sample folder first
source
1280×720 · 25.000 fps · 800 fr
/srv/tw/video/abstract.mp4
analyse / open
recent projects
samples
/srv/tw/corpus
analyse folder
all
none
sample space
tone ↔ · grit ↕
dots = samples (size: activity). rings = what each layer is playing at the playhead. the image moves a target through this space.
frame 180
lanes
brightness
hue red↔cyan
motion
detail / edges
pitch
lowpass
+ lane
drag = mark · wheel = zoom · shift+wheel = scroll
fit
00:00.00
00:05.00
00:10.00
00:15.00
00:20.00
00:25.00
00:30.00
no render yet — press render
regions 0 placed · drag on the timeline to mark a part to redo · enter regenerate · click a region to select it
mapping
layers
engine
fx
export
preset
— load preset —
mutate
random
clear
pitch
scan speed
grain size
density
spray
reverse
level
pan
width
lowpass
highpass
resonance
drive
chorus
pulse
echo
reverb
space size
match · tone
match · grit
match · activity
match · height
brightness
contrast
flicker
balance L↔R ±
hue red↔cyan
hue yellow↔blue
colour shift
saturation
warmth ±
colourfulness
detail / edges
sharpness
tonal entropy
motion
motion ← → ±
motion ↓ ↑ ±
zoom out↔in ±
rotation ±
motion pos x ±
motion spread
motion coherence
frame change
cut impulse
drag a cell up/down · double-click toggles · right-click clears
sensitivity
1.00
base values
value with zero modulation
pitch 0.50
scan speed 0.25
grain size 0.50
density 0.50
spray 0.12
reverse 0.00
level 0.80
pan 0.50
width 0.50
lowpass 0.85
highpass 0.06
resonance 0.08
drive 0.00
chorus 0.08
pulse 0.00
echo 0.06
reverb 0.30
space size 0.40
match · tone 0.50
match · grit 0.50
match · activ… 0.50
match · height 0.50

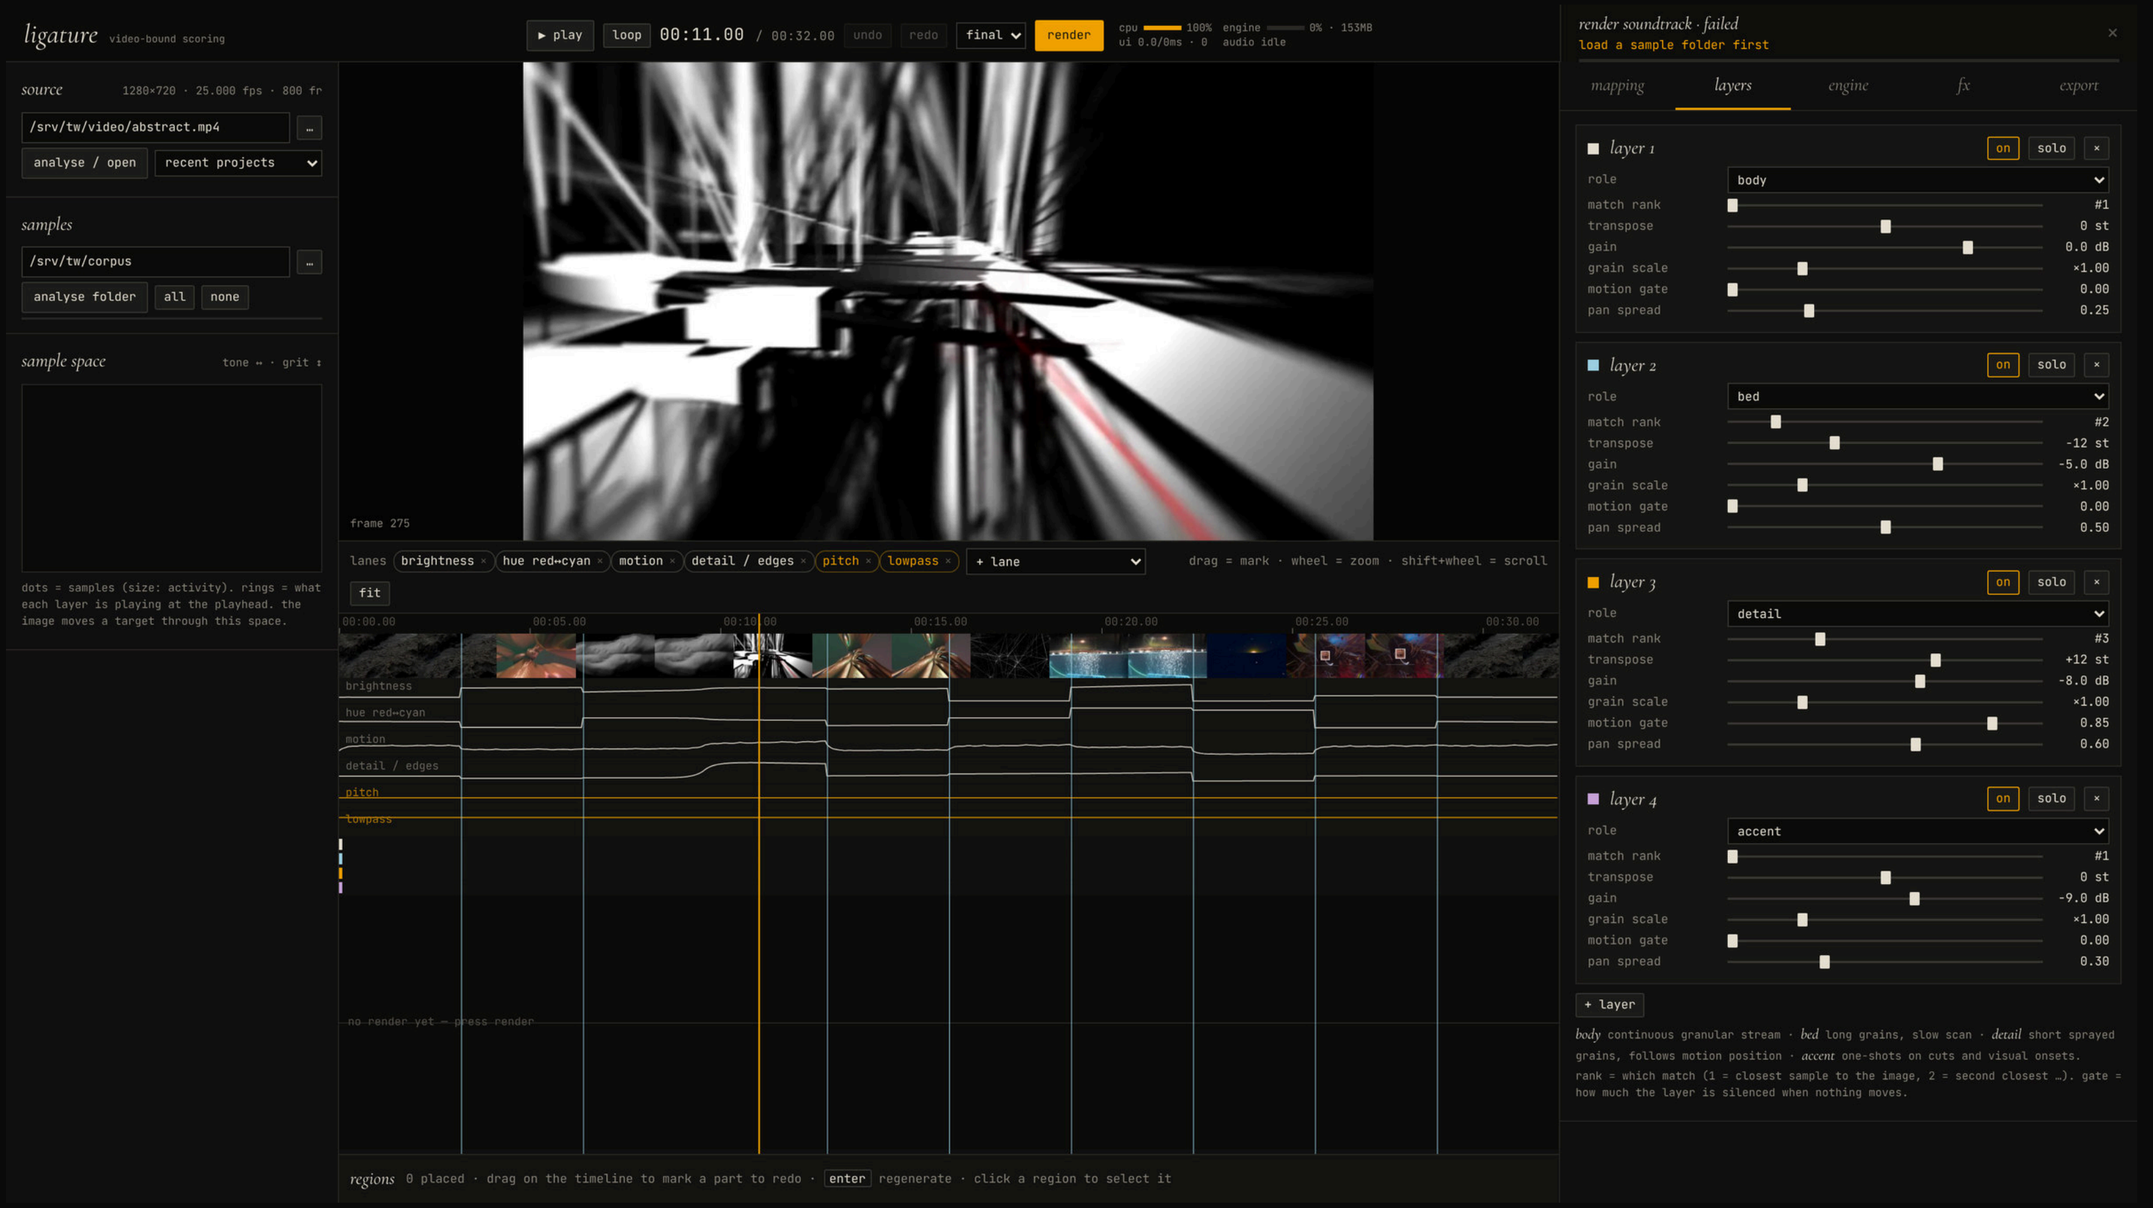
T05

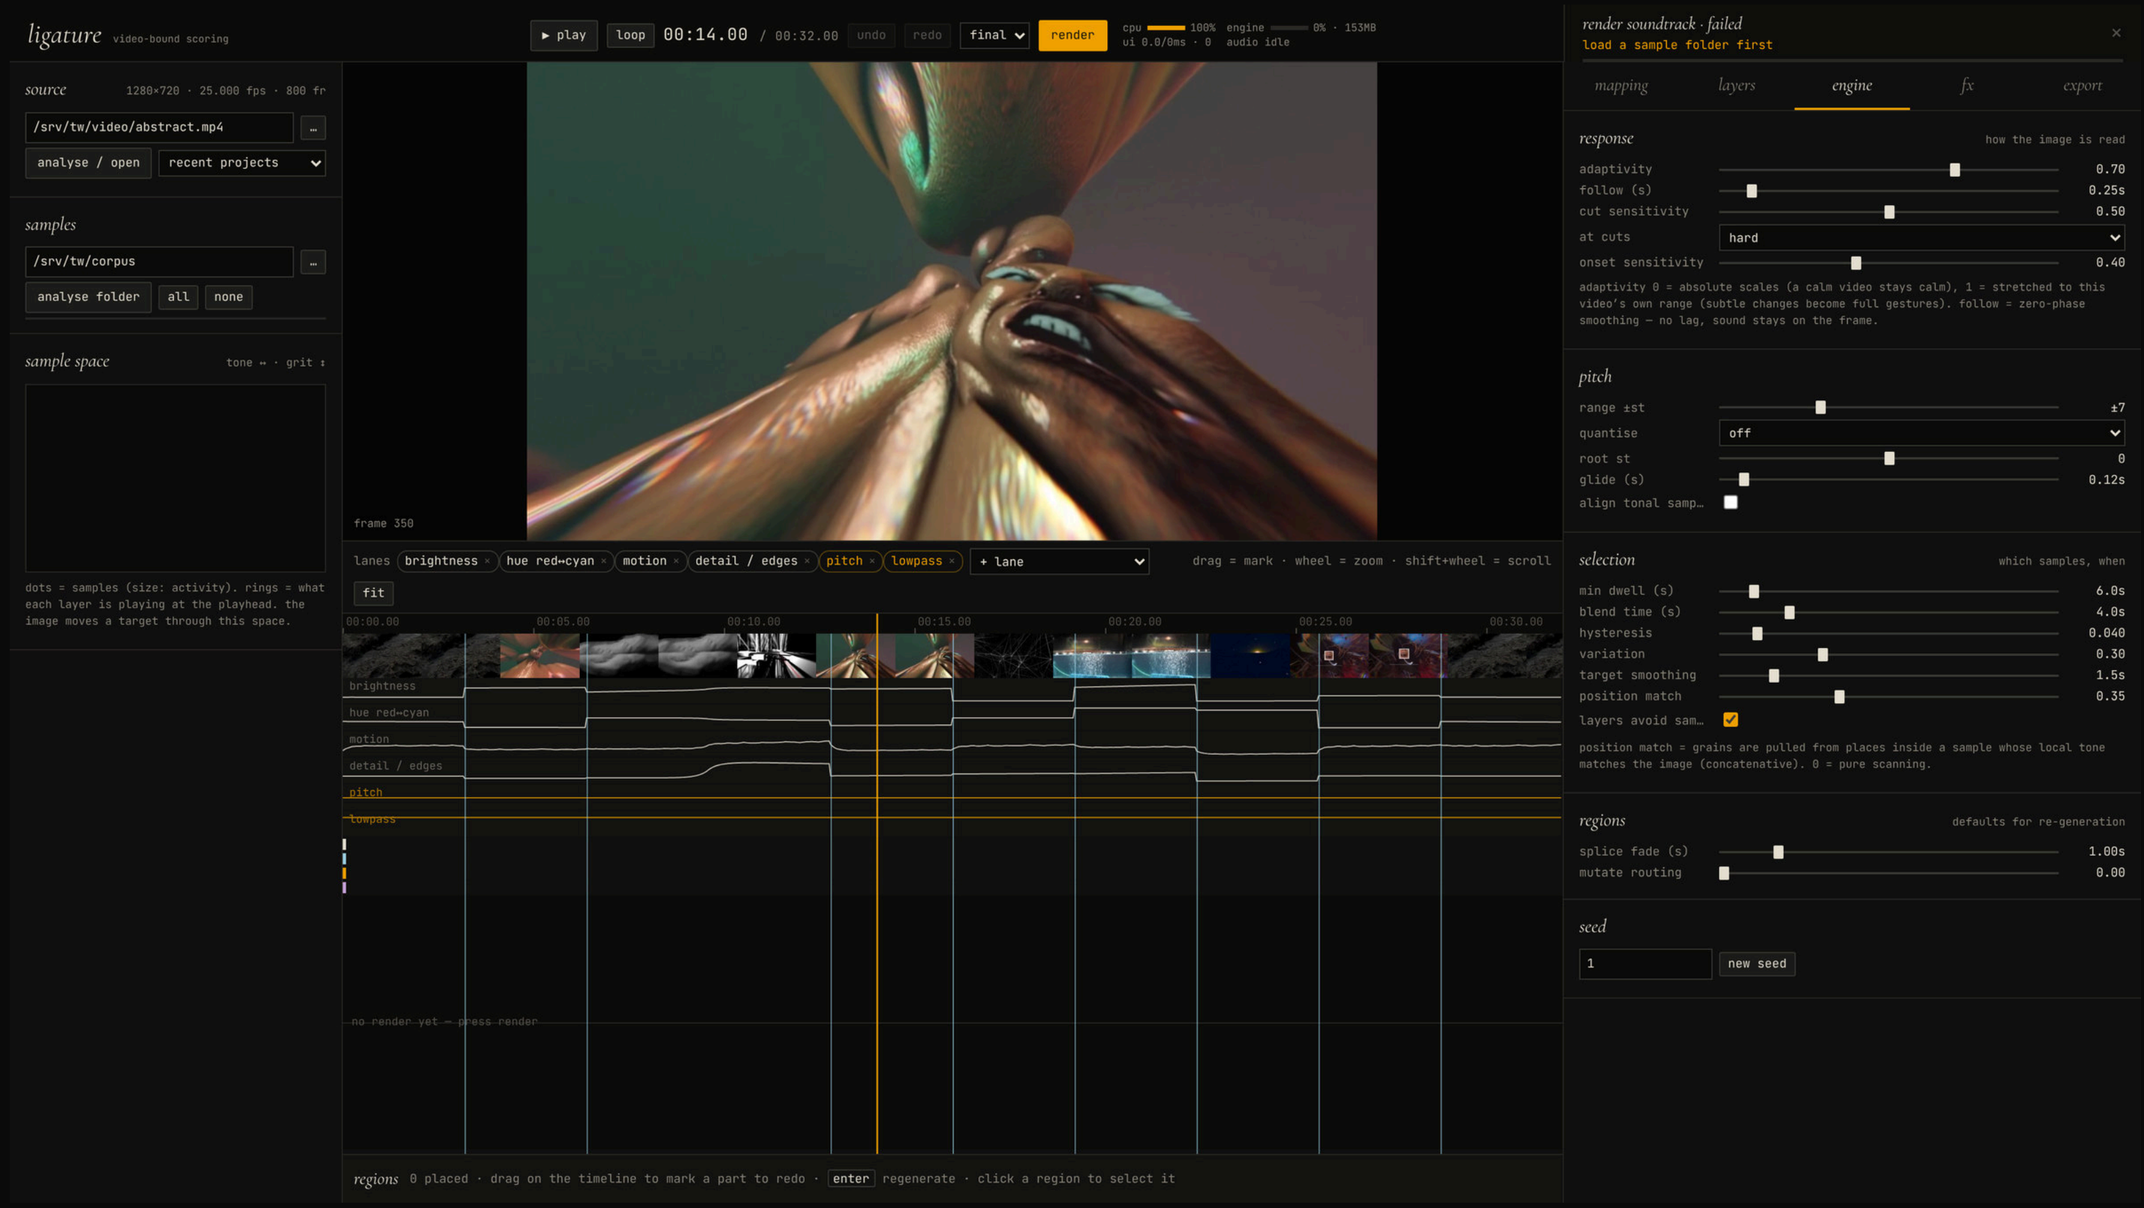
ligature

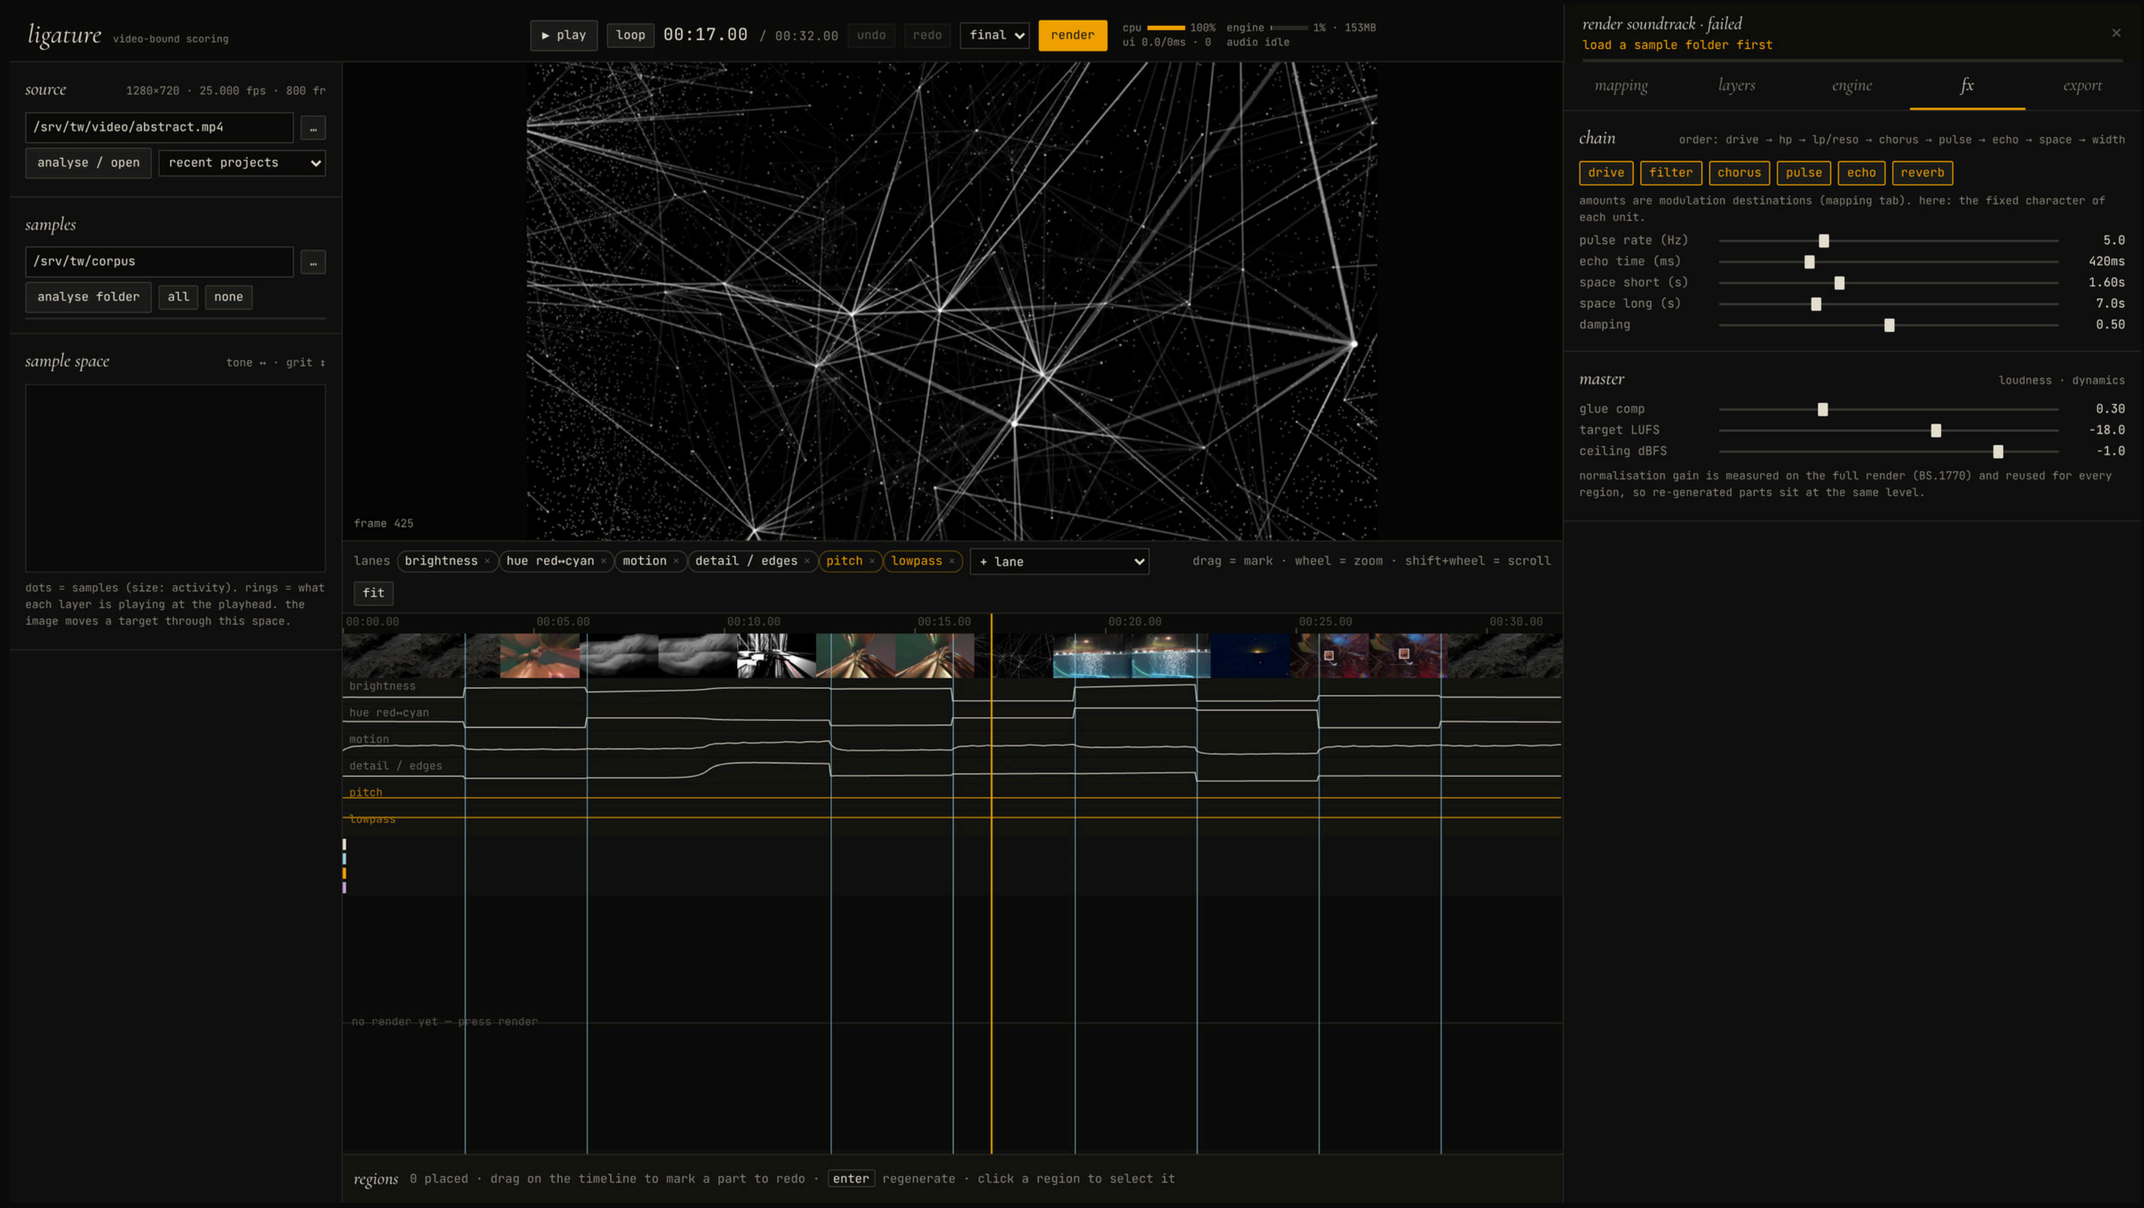
2026

FIELD	Picture-bound sound — audiovisual scoring
TECHNOLOGY	Python engine (numpy/scipy), HTML/Web Audio UI, imageio-ffmpeg
SOURCE	261006_ligature.zip

Reads a video frame by frame and builds a soundtrack from a sample folder. Image descriptors — brightness, hue axes, motion and its direction, edge density, cut impulses — are routed through a modulation matrix to granular parameters: which sample sounds, where in it grains are read, pitch, grain size, scan speed, filtering, space and stereo position. Picture and sound are bound sample-accurately; regions can be re-rendered locally.


ligature
video-bound scoring
▶ play
loop
00:11.00 / 00:32.00
undo
redo
final
render
cpu 100% engine 0% · 153MB
ui 0.0/0ms · 0 audio idle
render soundtrack · failed
load a sample folder first
mapping
layers
engine
fx
export
source
1280×720 · 25.000 fps · 800 fr
/srv/tw/video/abstract.mp4
analyse / open
recent projects
samples
/srv/tw/corpus
analyse folder
all
none
sample space
tone ↔ · grit ↕
dots = samples (size: activity). rings = what each layer is playing at the playhead. the image moves a target through this space.
frame 275
lanes
brightness
hue red↔cyan
motion
detail / edges
pitch
lowpass
+ lane
drag = mark · wheel = zoom · shift+wheel = scroll
fit
00:00.00
00:05.00
00:10.00
00:15.00
00:20.00
00:25.00
00:30.00
brightness
hue red↔cyan
motion
detail / edges
pitch
lowpass
no render yet — press render
regions
0 placed · drag on the timeline to mark a part to redo · enter regenerate · click a region to select it
layer 1
on
solo
role
body
match rank #1
transpose 0 st
gain 0.0 dB
grain scale ×1.00
motion gate 0.00
pan spread 0.25
layer 2
on
solo
role
bed
match rank #2
transpose -12 st
gain -5.0 dB
grain scale ×1.00
motion gate 0.00
pan spread 0.50
layer 3
on
solo
role
detail
match rank #3
transpose +12 st
gain -8.0 dB
grain scale ×1.00
motion gate 0.85
pan spread 0.60
layer 4
on
solo
role
accent
match rank #1
transpose 0 st
gain -9.0 dB
grain scale ×1.00
motion gate 0.00
pan spread 0.30
+ layer
body continuous granular stream · bed long grains, slow scan · detail short sprayed grains, follows motion position · accent one-shots on cuts and visual onsets.
rank = which match (1 = closest sample to the image, 2 = second closest …). gate = how much the layer is silenced when nothing moves.


ligature
video-bound scoring
▶ play
loop
00:14.00 / 00:32.00
undo
redo
final
render
cpu 100% engine 0% · 153MB
ui 0.0/0ms · 0 audio idle
render soundtrack · failed
load a sample folder first
mapping
layers
engine
fx
export
source
1280×720 · 25.000 fps · 800 fr
/srv/tw/video/abstract.mp4
analyse / open
recent projects
samples
/srv/tw/corpus
analyse folder
all
none
sample space
tone ↔ · grit ↕
dots = samples (size: activity). rings = what each layer is playing at the playhead. the image moves a target through this space.
frame 350
lanes
brightness
hue red↔cyan
motion
detail / edges
pitch
lowpass
+ lane
drag = mark · wheel = zoom · shift+wheel = scroll
fit
00:00.00
00:05.00
00:10.00
00:15.00
00:20.00
00:25.00
00:30.00
no render yet — press render
regions
0 placed · drag on the timeline to mark a part to redo · enter regenerate · click a region to select it
response
how the image is read
adaptivity 0.70
follow (s) 0.25s
cut sensitivity 0.50
at cuts hard
onset sensitivity 0.40
adaptivity 0 = absolute scales (a calm video stays calm), 1 = stretched to this video's own range (subtle changes become full gestures). follow = zero-phase smoothing — no lag, sound stays on the frame.
pitch
range ±st ±7
quantise off
root st 0
glide (s) 0.12s
align tonal samp…
selection
which samples, when
min dwell (s) 6.0s
blend time (s) 4.0s
hysteresis 0.040
variation 0.30
target smoothing 1.5s
position match 0.35
layers avoid sam…
position match = grains are pulled from places inside a sample whose local tone matches the image (concatenative). 0 = pure scanning.
regions
defaults for re-generation
splice fade (s) 1.00s
mutate routing 0.00
seed
1
new seed


ligature
video-bound scoring
▶ play
loop
00:17.00 / 00:32.00
undo
redo
final
render
cpu 100% engine 1% · 153MB
ui 0.0/0ms · 0 audio idle
source
1280×720 · 25.000 fps · 800 fr
/srv/tw/video/abstract.mp4
…
analyse / open
recent projects
samples
/srv/tw/corpus
…
analyse folder
all
none
sample space
tone ↔ · grit ↕
dots = samples (size: activity). rings = what each layer is playing at the playhead. the image moves a target through this space.
frame 425
lanes
brightness
hue red↔cyan
motion
detail / edges
pitch
lowpass
+ lane
drag = mark · wheel = zoom · shift+wheel = scroll
fit
00:00.00
00:05.00
00:10.00
00:15.00
00:20.00
00:25.00
00:30.00
brightness
hue red↔cyan
motion
detail / edges
pitch
lowpass
no render yet — press render
regions
0 placed · drag on the timeline to mark a part to redo · enter regenerate · click a region to select it
render soundtrack · failed
load a sample folder first
mapping
layers
engine
fx
export
chain
order: drive → hp → lp/reso → chorus → pulse → echo → space → width
drive
filter
chorus
pulse
echo
reverb
amounts are modulation destinations (mapping tab). here: the fixed character of each unit.
pulse rate (Hz) 5.0
echo time (ms) 420ms
space short (s) 1.60s
space long (s) 7.0s
damping 0.50
master
loudness · dynamics
glue comp 0.30
target LUFS -18.0
ceiling dBFS -1.0
normalisation gain is measured on the full render (BS.1770) and reused for every region, so re-generated parts sit at the same level.



sortcut

2026

PICTURE-BOUND SOUND

01 · chromatic drift · sort strip

T06

T06

sortcut

2026

FIELD	Picture-bound sound — audiovisual scoring
VERSION	1.0
TECHNOLOGY	Python (numpy, OpenCV, torch/CUDA: CLIP, YOLO, RAFT, Depth Anything, MediaPipe), browser UI, NVDEC/NVENC, OpenTimelineIO
SOURCE	260929_sortcut.zip

A generative montage system for surveillance footage. Every frame of every camera is described once — colour, light, motion and optical flow, presence, composition, CLIP semantics, detections, faces, pose, depth and real clock time — and the material is then cut by sorting rather than by chance: 74 presets in families such as perceptual, continuity, semantic, surveillance, people, faces, image, depth, time, structure and audio. Continuity constraints, curve readings of every ranking and a use-once rule turn sorting into editing; each render is reproducible from its sidecar (CSV, OTIO, JSON) and can be streamed as OSC for sonification.

sortcut /srv/tw/cctv

BROWSE OPEN

CPU ONLY NVDEC NVENC

SOURCES

6/6 · 1440 F

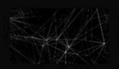
INGEST RESCAN ALL NONE



01 cam01
1280×720 · 240 f · 1.00 fps
time: index



02 cam02
1280×720 · 240 f · 1.00 fps
time: index



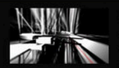
03 cam03
1280×720 · 240 f · 1.00 fps
time: index



04 cam04
1280×720 · 240 f · 1.00 fps
time: index



05 cam05
1280×720 · 240 f · 1.00 fps
time: index



06 cam06
1280×720 · 240 f · 1.00 fps
time: index

INGEST OPTIONS +

ANALYSIS

7/12 COMPLETE

- color 6/6
- motion 6/6
- background 6/6
- structure 6/6
- flow 6/6
- embed 6/6
- detect n/a
- face n/a
- pose n/a
- depth n/a
- ocr n/a
- rarity 6/6

ANALYZER OPTIONS +

redo

ANALYZE

SEQUENCE EXPLORE RENDERS

PRESETS

74

PERCEPTUAL

01 chromatic drift
02 chromatic drift · grains
03 lab path
04 synthetic day
05 synthetic night
06 motion swell
07 motion breath
08 stillness

CONTINUITY

09 impossible camera
10 impossible camera · composition
11 match cut
12 match cut · semantic
13 anti-continuity
14 drift apart
15 motion relay
16 motion sweep
17 wind rose

SEMANTIC

18 prompt: empty → crowd
19 prompt: weather
20 prompt: waiting
21 unsupervised axis
22 principal axis
23 typology
24 object chain

SURVEILLANCE

25 · population
26 · ghost
27 rarity
28 rarity · only anomalies
29 panopticon
30 panopticon · long stare
31 absence
32 only absence

PEOPLE

33 · empty → crowded
34 · crowd swell
35 · approach
36 · arms rising
37 · falling
38 · pose rhyme
39 · stride

FACES

40 · eye line
41 · face lock · head turn
42 · face lock · morph
43 · face scale
44 · scream
45 · smile
46 · eyes closing
47 · gaze sweep
48 · expression rhyme

IMAGE

49 void → noise
50 texture scale
51 focus pull
52 rotating world
53 visual weight
54 monochrome → polychrome
55 cold → warm
56 symmetry

DEPTH

57 · tunnel → wall
58 · near → far
59 · depth rhyme

TIME

60 one day

STRUCTURAL

61 hue × light hilbert
62 semantic spiral
63 converge
64 fractal brightness
65 light then colour
66 · people then place
67 verticals
68 wall · per cam brightness
69 canon · 4 screens

AUDIO

70 audio: loudness → light
71 audio: brightness → motion
72 audio: stutter

BASELINE

73 original
74 chronological merge

the whole piece becomes one colour gradient across all cams

SEQUENCE history

PARAMETERS RENDER JOBS

METHOD

chromatic drift

method: sort by feature

rank all units by one feature, then read the ranking along a curve

key: hue_dom — dominant hue

desc: on off

curve: linear

cycles: 3.0

circ start: auto fixed

circ deg: 0

gray mode: end start drop keep

chroma min: 6.0

CONTINUITY

once: on off

tight: 0.90

cont space: color+L,a,b

UNIT

unit: frame grain event scene

grain len: 8

grain jitter: 0.00

grain dir: forward reverse pingpong random

event signal: motion bgdiff persons fg_area objects

event q: 0.75

event min: 2

event max: 60

event pad: 1

event merge: 3

scene q: 0.97

SEQUENCE

sortcut /srv/tw/cctv BROWSE OPEN CPU ONLY NVDEC NVENC

SOURCES 6/6 · 1440 F
INGEST RESCAN ALL NONE

01 cam01 ■
1280×720 · 240 f · 1.00 fps
time: index

02 cam02 ■
1280×720 · 240 f · 1.00 fps
time: index

03 cam03 ■
1280×720 · 240 f · 1.00 fps
time: index

04 cam04 ■
1280×720 · 240 f · 1.00 fps
time: index

05 cam05 ■
1280×720 · 240 f · 1.00 fps
time: index

06 cam06 ■
1280×720 · 240 f · 1.00 fps
time: index

INGEST OPTIONS +

ANALYSIS 7/12 COMPLETE
- color 6/6
- motion 6/6
- background 6/6
- structure 6/6
- flow 6/6
- embed 6/6
- detect n/a
- face n/a
- pose n/a
- depth n/a
- ocr n/a
- rarity 6/6

ANALYZER OPTIONS +

redo ANALYZE

CHROMATIC DRIFT
1440 UNITS → 1440 F · 48S · 1440 CUTS · STEP 0.014 · 100% UNIQUE · 0.14S

cam06 · f8 · key 22.302

PAUSE ⟲ START ×.25 ×.5 ×1 ×2 ×4 ■ loop 1320 / 1440 0:44.0 / 0:48.0

SOURCE COLOUR PER FRAME — SORT KEY | CUTS ▬ CAMERA

0:00.0 0:00.2 0:00.4 0:00.6 0:00.8 0:01.0 0:01.2 0:01.4 0:01.6 0:01.8 0:02.0
0:02.2 0:02.4 0:02.6 0:02.8 0:03.0 0:03.2 0:03.4 0:03.6 0:03.8 0:04.0 0:04.2
0:04.4 0:04.6 0:04.8 0:05.0 0:05.2 0:05.4 0:05.6 0:05.8 0:06.0 0:06.2 0:06.4
0:06.6 0:06.8 0:07.0 0:07.2 0:07.4 0:07.6 0:07.8 0:08.0 0:08.2 0:08.4 0:08.6

PARAMETERS RENDER JOBS

METHOD chromatic drift

◆ method: sort by feature
rank all units by one feature, then read the ranking along a curve

⇅ key: hue_dom — dominant hue

↓ desc: on off

∫ curve: linear

∿ cycles: 3.0

○ circ start: auto fixed

° circ deg: 0

▒ gray mode: end start drop keep

◌ chroma min: 6.0

CONTINUITY

1× once: on off

⊶ tight: 0.90

≈ cont space: color+L,a,b

UNIT

■ unit: frame grain event scene

■■■ grain len: 8

⁖ grain jitter: 0.00

⇄ grain dir: forward reverse pingpong random

⌁ event signal: motion bgdiff persons fg_area objects

≥ event q: 0.75

⌊ event min: 2

⌈ event max: 60

⟨⟩ event pad: 1

⌒ event merge: 3

§ scene q: 0.97

48s · 1440 cuts SEQUENCE

sortcut /srv/tw/cctv BROWSE OPEN CPU ONLY NVDEC NVENC

SOURCES 6/6 · 1440 F
INGEST RESCAN ALL NONE

01 cam01 ■
1280×720 · 240 f · 1.00 fps
time: index

02 cam02 ■
1280×720 · 240 f · 1.00 fps
time: index

03 cam03 ■
1280×720 · 240 f · 1.00 fps
time: index

04 cam04 ■
1280×720 · 240 f · 1.00 fps
time: index

05 cam05 ■
1280×720 · 240 f · 1.00 fps
time: index

06 cam06 ■
1280×720 · 240 f · 1.00 fps
time: index

INGEST OPTIONS +

ANALYSIS 7/12 COMPLETE
- color 6/6
- motion 6/6
- background 6/6
- structure 6/6
- flow 6/6
- embed 6/6
- detect n/a
- face n/a
- pose n/a
- depth n/a
- ocr n/a
- rarity 6/6

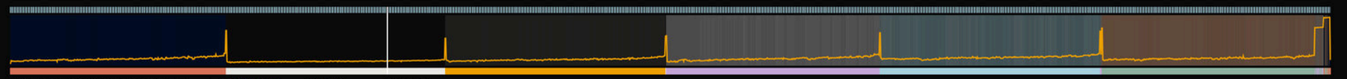
ANALYZER OPTIONS +

redo ANALYZE

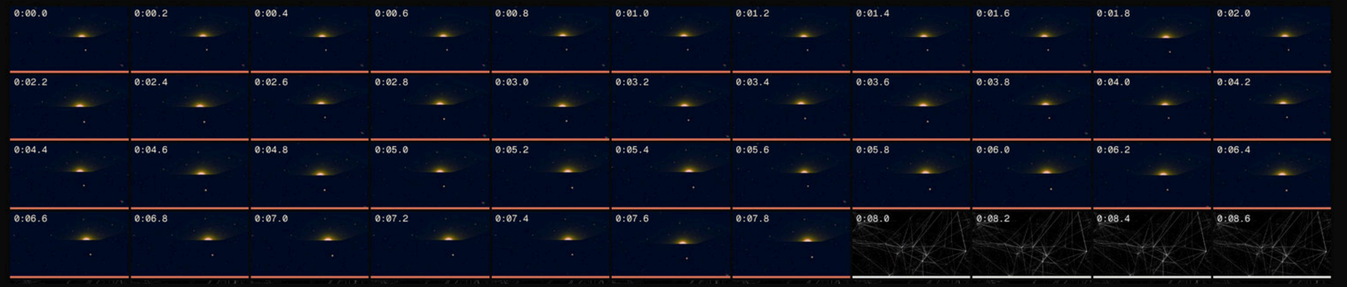
IMPOSSIBLE CAMERA
1440 UNITS → 1440 F · 48S · 1440 CUTS · STEP 0.022 · 100% UNIQUE · 1S



PAUSE ⟲ START ×.25 ×.5 ×1 ×2 ×4 ■ loop 412 / 1440 0:13.7 / 0:48.0



SOURCE COLOUR PER FRAME — SORT KEY | CUTS ▬ CAMERA



PARAMETERS RENDER JOBS

METHOD impossible camera
method: shortest path (impossib
one continuous path through feature space visiting every unit once — cuts that feel like no cut
space: embed+color
min jump: 30
cam bias: 0.40
start: darkest
opt seconds: 20

CONTINUITY
once: on off
tight: 0.60
cont space: color+layout

UNIT
unit: frame grain event scene
grain len: 8
grain jitter: 0.00
grain dir: forward reverse pingpong random
event signal: motion bgdiff persons fg_area objects
event q: 0.75
event min: 2
event max: 60
event pad: 1
event merge: 3
scene q: 0.97

FILTER +

TIME
fps: 30

48s · 1440 cuts SEQUENCE


scorecut 1.3
abstract.mp4 1280×720 25.000 fps · cfr 800 frames 0:32 flow · dis
random all
undo
analyze done
SOURCE
/srv/tw/video/abstract.mp4
pick
audio file or folder
files
folder
srv/tw/corpus/
analyse
cached after the first pass
PRESET
revert
CHANNEL clean tuner resume dead-air
OPTIC focus luminance grain depth exposure
CHROMA fifths common warmth prism
KINETIC stillness drift pressure tape
SPECTRAL lattice shift entropy mirror spatial
VOID night infrared sodium
STRANGE height camera novelty total
MATCH corpus sorted
dead air + relay tick at every cut; image grain opens the carrier hiss
save as …
save
BINDINGS
preset depth 1.0
unleash off on
file dice roll
position dice roll
SCENES
SELECTION
file per scene
channel match sorted sequence random single
position in file
broadcast resume head match song
variety 0.40
distinct neighbours off on
avoid silence off on
sort axis
bright energy tone texture mass
luma ↔ centroid 1.00
motion ↔ onsets 0.80
saturation ↔ tonality 0.60
00:00:06:15
frame 000165
scene 3/10 · cam 3
l2.wav
@ 31.60s · -0.7 dB
proxy · silent
play
audition
proxy
loop
flow
files
positions
00:00:06:15 / 00:00:32:00
3 1 3 5 6 4 5 2 6 4
sensor noise → noise floor
MAPPINGS · IMAGE → SOUND
+ map
sensor noise → noise floor
0 … 0.35 × 1 route
frame scene
lin exp log s st4 st8 gate
smooth 120 ms
MOTION LAYER · OPTICAL FLOW
+ map
FLOW off imprint carve tide doppler scrub flutter swarm rotor impulse gaze vertical wake full
custom layer
motion layer off on
layer depth 1.0
imprint
off spotlight carve boost
imprint depth 18 dB
imprint pan 0.60
imprint curve 1.0
imprint attack 6 ms
imprint release 240 ms
bottom row 70 Hz
top row 12000 Hz
flutter shape sine pulse
flutter restarts at cut off on
LAYER MAPPINGS – ADDED ON TOP OF ANY PRESET
pick a flow preset above or + map – flow sources are in the 'flow' group
CORPUS
6 files
file dur key LUFS centr scn
l0.wav 0:30 G#m -22.5 426 1
l1.wav 0:36 A#m -23.5 107 1
l2.wav 0:42 C# -22.2 228 2
l3.wav 0:48 C# -19.3 359 2
l4.wav 0:54 A#m -21.2 385 2
l5.wav 1:00 C# -19.1 317 2
render
cancel
800 frames · 6 files

T07

scorecut

2026

FIELD	Picture-bound sound — audiovisual scoring
VERSION	1.3
TECHNOLOGY	Python (numpy, scipy, OpenCV optical flow), embedded browser UI, ffmpeg
SOURCE	scorecut_1.3.zip

Companion to sortcut (T06): sortcut makes the cut, scorecut makes it sound. The cut video is decomposed into scenes and cameras; each camera becomes a ‘station’ of an audio corpus, switched at the exact frame of the cut, while image features drive a spectral and spatial treatment of the sound. Thirty presets in eight families; cut-aware smoothing guarantees that every change lands on its frame. Output: high-resolution audio, the stream-copied picture with new sound, and a cue list.

scorecut 1.3 — abstract.mp4 — 1280×720 — 25.000 fps · cfr — 800 frames — 0:32 — flow · dis — random all — undo — analyze done

SOURCE

/srv/tw/video/abstract.mp4 [pick]

audio file or folder [files] [folder]

srv/tw/corpus/ ×

[analyse] cached after the first pass

PRESET revert

CHANNEL clean tuner resume dead-air
OPTIC focus luminance grain depth exposure
CHROMA fifths common warmth prism
KINETIC stillness drift pressure tape
SPECTRAL lattice shift entropy mirror spatial
VOID night infrared sodium
STRANGE height camera novelty total
MATCH corpus sorted

each colour a transposition, stepped per scene; saturation opens the stereo field

save as … [save]

BINDINGS

preset depth 1.0 ×
unleash off on
file dice roll ×
position dice roll ×

SCENES

SELECTION

file per scene
channel match sorted sequence random single
position in file
broadcast resume head match song
variety 0.40
distinct neighbours off on
avoid silence off on
sort axis
bright energy tone texture mass
luma ↔ centroid 1.00
motion ↔ onsets 0.80
saturation ↔ tonality 0.60

00:00:13:12
frame 000337

scene 5/10 · cam 5
l5.wav
@ 26.64s · +0.9 dB

proxy · silent

▶ play ◉ audition ◌ proxy ↻ loop → flow files positions 00:00:13:12 / 00:00:32:00

3 1 3 5 6 4 5 2 6 4

hue → pitch

saturation → width

MAPPINGS · IMAGE → SOUND + map

hue → pitch ×
-6 … 6 st × 1 route frame scene
lin exp log s st4 st8 gate
smooth 0 ms

saturation → width ×
0.3 … 1.8 × 1 route frame scene
lin exp log s st4 st8 gate
smooth 80 ms

MOTION LAYER · OPTICAL FLOW + map

FLOW off imprint carve tide doppler scrub flutter swarm rotor impulse gaze vertical wake full

custom layer

motion layer off on
layer depth 1.0 ×
imprint
off spotlight carve boost
imprint depth 18 dB
imprint pan 0.60
imprint curve 1.0
imprint attack 6 ms
imprint release 240 ms
bottom row 70 Hz
top row 12000 Hz
flutter shape sine pulse
flutter restarts at cut off on

LAYER MAPPINGS — ADDED ON TOP OF ANY PRESET

pick a flow preset above or + map — flow sources are in the 'flow' group

CORPUS 6 files

file	dur	key	LUFS	centr	scn
l0.wav	0:30	G#m	-22.5	426	1
l1.wav	0:36	A#m	-23.5	107	1
l2.wav	0:42	C#	-22.2	228	2
l3.wav	0:48	C#	-19.3	359	2
l4.wav	0:54	A#m	-21.2	385	2
l5.wav	1:00	C#	-19.1	317	2

[render] [cancel] 800 frames · 6 files

scorecut 1.3 — abstract.mp4 — 1280×720 — 25.000 fps · cfr — 800 frames — 0:32 — flow · dis — random all — undo — analyze done

SOURCE

/srv/tw/video/abstract.mp4 — pick

audio file or folder — files — folder

srv/tw/corpus/
srv/tw/samples/s00.wav/
srv/tw/samples/s01.wav/
srv/tw/samples/s02.wav/
srv/tw/samples/s03.wav/
srv/tw/samples/s04.wav/
srv/tw/samples/s05.wav/

analyse — cached after the first pass

PRESET — revert

CHANNEL clean tuner resume dead-air
OPTIC focus luminance grain depth exposure
CHROMA fifths common warmth prism
KINETIC stillness drift pressure tape
SPECTRAL lattice shift entropy mirror spatial
VOID night infrared sodium
STRANGE height camera novelty total
MATCH corpus sorted

horizontal architecture sustains, vertical structure clicks

save as … — save

BINDINGS

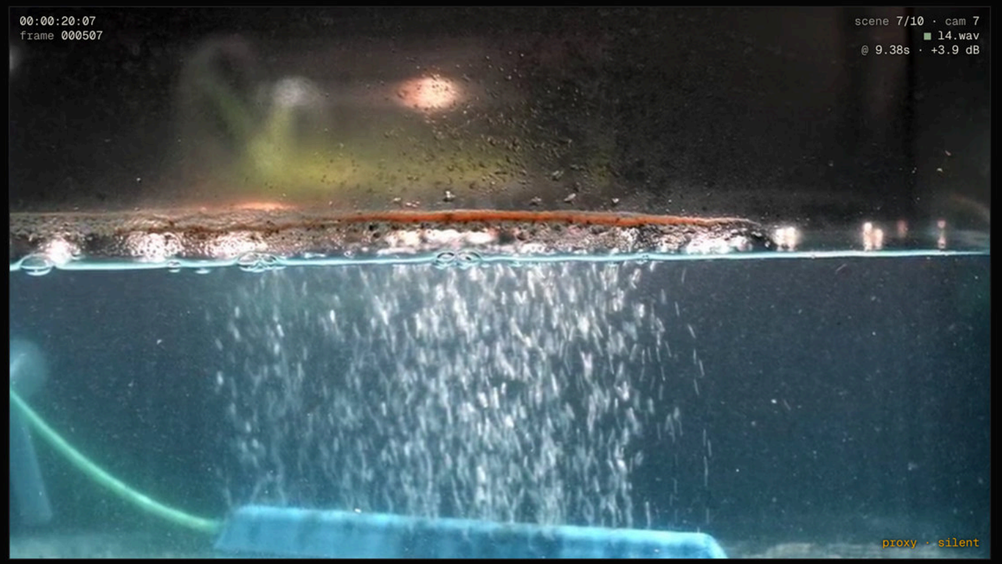
preset depth 1.0
unleash off on
file dice roll
position dice roll

SCENES

SELECTION

file per scene
channel match sorted sequence random single
position in file
broadcast resume head match song
variety 0.40
distinct neighbours off on
avoid silence off on
sort axis
bright energy tone texture mass
luma ↔ centroid 1.00

render — cancel — 800 frames · 6 files



▶ play — audition — proxy — loop — flow — files — positions — 00:00:20:07 / 00:00:32:00



MAPPINGS · IMAGE → SOUND — + map

orientation → harm ↔ perc
-0.9 … 0.9 × 1 route — frame scene
lin exp log s st4 st8 gate
smooth — 100 ms

detail → resolution loss
0 … 0.8 × 1 route — frame scene
lin exp log s st4 st8 gate
smooth — 100 ms

entropy → diffusion
0 … 0.5 × 1 route — frame scene
lin exp log s st4 st8 gate
smooth — 100 ms

MOTION LAYER · OPTICAL FLOW — + map

FLOW off imprint carve tide doppler scrub flutter swarm rotor impulse gaze vertical wake full

custom layer

motion layer off on
layer depth 1.0
imprint
off spotlight carve boost
imprint depth 18 dB
imprint pan 0.60
imprint curve 1.0
imprint attack 6 ms
imprint release 240 ms
bottom row 70 Hz
top row 12000 Hz
flutter shape sine pulse
flutter restarts at cut off on

LAYER MAPPINGS — ADDED ON TOP OF ANY PRESET

pick a flow preset above or + map — flow sources are in the 'flow' group

CORPUS — 6 files

file	dur	key	LUFS	centr	scn
l0.wav	0:30	G#m	-22.5	426	1
l1.wav	0:36	A#m	-23.5	107	1
l2.wav	0:42	C#	-22.2	228	2
l3.wav	0:48	C#	-19.3	359	2

scorecut 1.3 — abstract.mp4 — 1280×720 — 25.000 fps · cfr — 800 frames — 0:32 — flow · dis — random all — undo — analyze done

SOURCE

/srv/tw/video/abstract.mp4 — pick

audio file or folder — files — folder

srv/tw/corpus/ ×
srv/tw/samples/s00.wav/ ×
srv/tw/samples/s01.wav/ ×
srv/tw/samples/s02.wav/ ×
srv/tw/samples/s03.wav/ ×
srv/tw/samples/s04.wav/ ×

analyse — cached after the first pass

PRESET — revert

CHANNEL clean tuner resume dead-air
OPTIC focus luminance grain depth exposure
CHROMA fifths common warmth prism
KINETIC stillness drift pressure tape
SPECTRAL lattice shift entropy mirror spatial
VOID night infrared sodium
STRANGE height camera novelty total
MATCH corpus sorted

horizontal architecture sustains, vertical structure clicks

save as … — save

BINDINGS

preset depth 1.0 ×
unleash off on
file dice roll ×
position dice roll ×

SCENES

SELECTION

file per scene
channel match sorted sequence random single
position in file
broadcast resume head match song
variety 0.40
distinct neighbours off on
avoid silence off on
sort axis
bright energy tone texture mass
luma ↔ centroid 1.00

render — cancel

positions re-rolled #471097



▶ play — audition — proxy — loop — flow — files — positions — 00:00:20:07 / 00:00:32:00



MAPPINGS · IMAGE → SOUND — + map

orientation → harm ↔ perc
-0.9 … 0.9 × 1 route — frame scene
lin exp log s st4 st8 gate
smooth — 100 ms

detail → resolution loss
0 … 0.8 × 1 route — frame scene
lin exp log s st4 st8 gate
smooth — 100 ms

entropy → diffusion
0 … 0.5 × 1 route — frame scene
lin exp log s st4 st8 gate
smooth — 100 ms

MOTION LAYER · OPTICAL FLOW — + map

FLOW off imprint carve tide doppler scrub flutter swarm rotor impulse gaze vertical wake full

custom layer

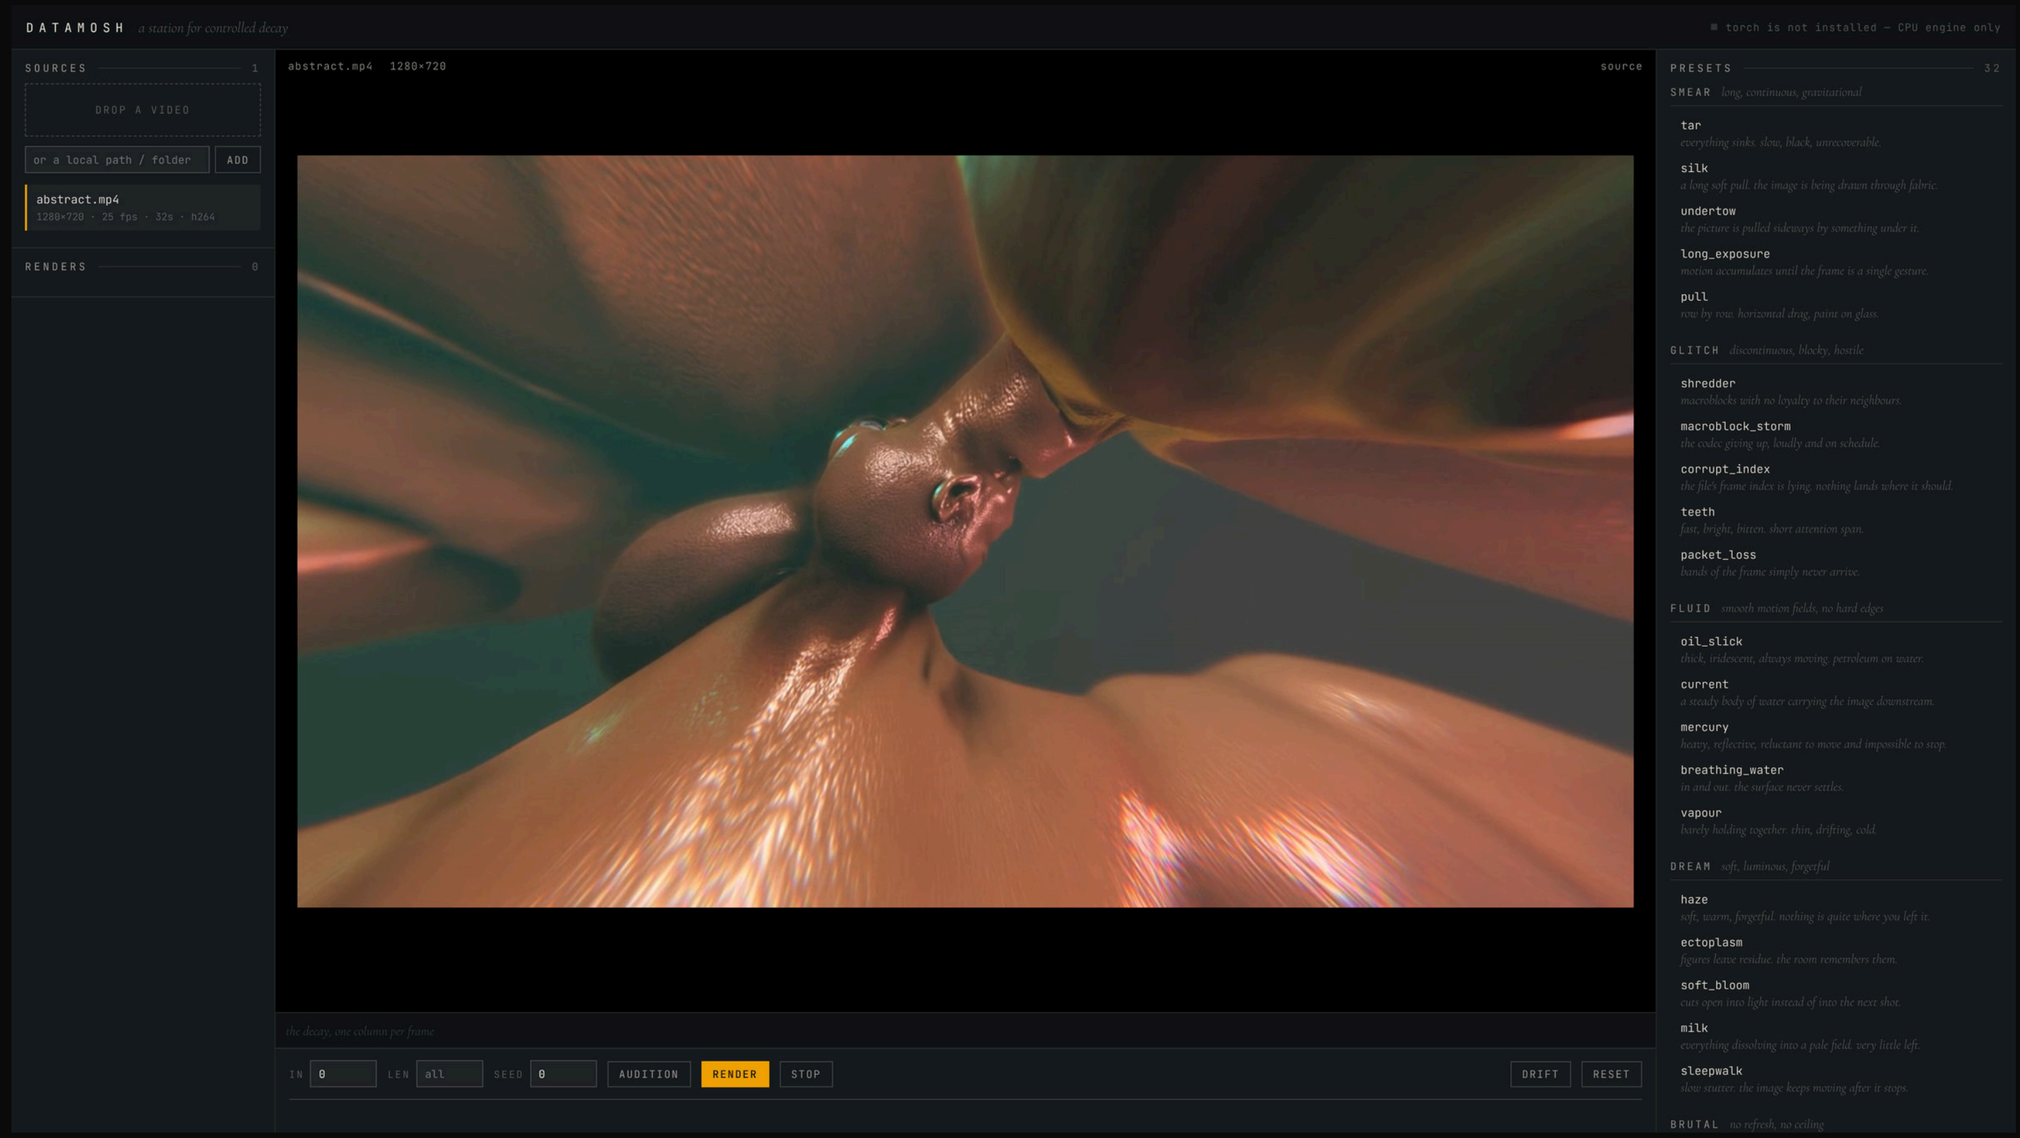
motion layer off on
layer depth 1.0 ×
imprint
off spotlight carve boost
imprint depth 18 dB
imprint pan 0.60
imprint curve 1.0
imprint attack 6 ms
imprint release 240 ms
bottom row 70 Hz
top row 12000 Hz
flutter shape sine pulse
flutter restarts at cut off on

LAYER MAPPINGS — ADDED ON TOP OF ANY PRESET

pick a flow preset above or + map — flow sources are in the 'flow' group

CORPUS — 6 files

file	dur	key	LUFS	centr	scn
l0.wav	0:30	G#m	-22.5	426	1
l1.wav	0:36	A#m	-23.5	107	1
l2.wav	0:42	C#	-22.2	228	2
l3.wav	0:48	C#	-19.3	359	2



datamosh station

2026

MOVING IMAGE

01 · source

T08

T08

datamosh station

2026

FIELD	Moving image — process, analysis, capture
VERSION	2
TECHNOLOGY	Python (numpy/OpenCV or CUDA/torch), browser UI, ffmpeg/xvid
SOURCE	260801_Datamosher2.zip

A standalone tool for controlled image decay. Two chainable engines: a flow engine in which measured motion is reshaped through 45 parameters and applied to a canvas that is never refreshed — repeatable at any resolution — and a bitstream engine that deletes I-frames and duplicates P-frames in real MPEG-4 data, leaving the damage to the decoder. Thirty-two pre-sets in five families (smear, glitch, fluid, dream, brutal); identical results on CUDA and CPU back-ends.

DATAMOSH *a station for controlled decay*

■ torch is not installed — CPU engine only

SOURCES 1

DROP A VIDEO

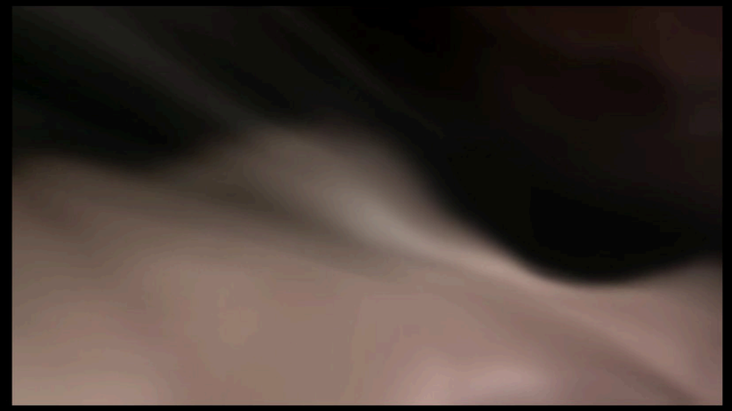
or a local path / folder | ADD

abstract.mp4
1280×720 · 25 fps · 32s · h264

RENDERS 1

fluid/mercury · audition done
99/100 · 37s folder view save

fluid/mercury · audition 640×360 · 99 frames render



IN 3.3 | LEN all | SEED 0 | AUDITION | RENDER | STOP | DRIFT | RESET

99 / 100 37s · done

PRESETS 32

SMEAR *long, continuous, gravitational*

tar
everything sinks. slow, black, unrecoverable.

silk
a long soft pull. the image is being drawn through fabric.

undertow
the picture is pulled sideways by something under it.

long_exposure
motion accumulates until the frame is a single gesture.

pull
row by row. horizontal drag, paint on glass.

GLITCH *discontinuous, blocky, hostile*

shredder
macroblocks with no loyalty to their neighbours.

macroblock_storm
the codec giving up, loudly and on schedule.

corrupt_index
the file's frame index is lying. nothing lands where it should.

teeth
fast, bright, bitten. short attention span.

packet_loss
bands of the frame simply never arrive.

FLUID *smooth motion fields, no hard edges*

oil_slick
thick, iridescent, always moving. petroleum on water.

current
a steady body of water carrying the image downstream.

mercury
heavy, reflective, reluctant to move and impossible to stop.

breathing_water
in and out. the surface never settles.

vapour
barely holding together. thin, drifting, cold.

DREAM *soft, luminous, forgetful*

haze
soft, warm, forgetful. nothing is quite where you left it.

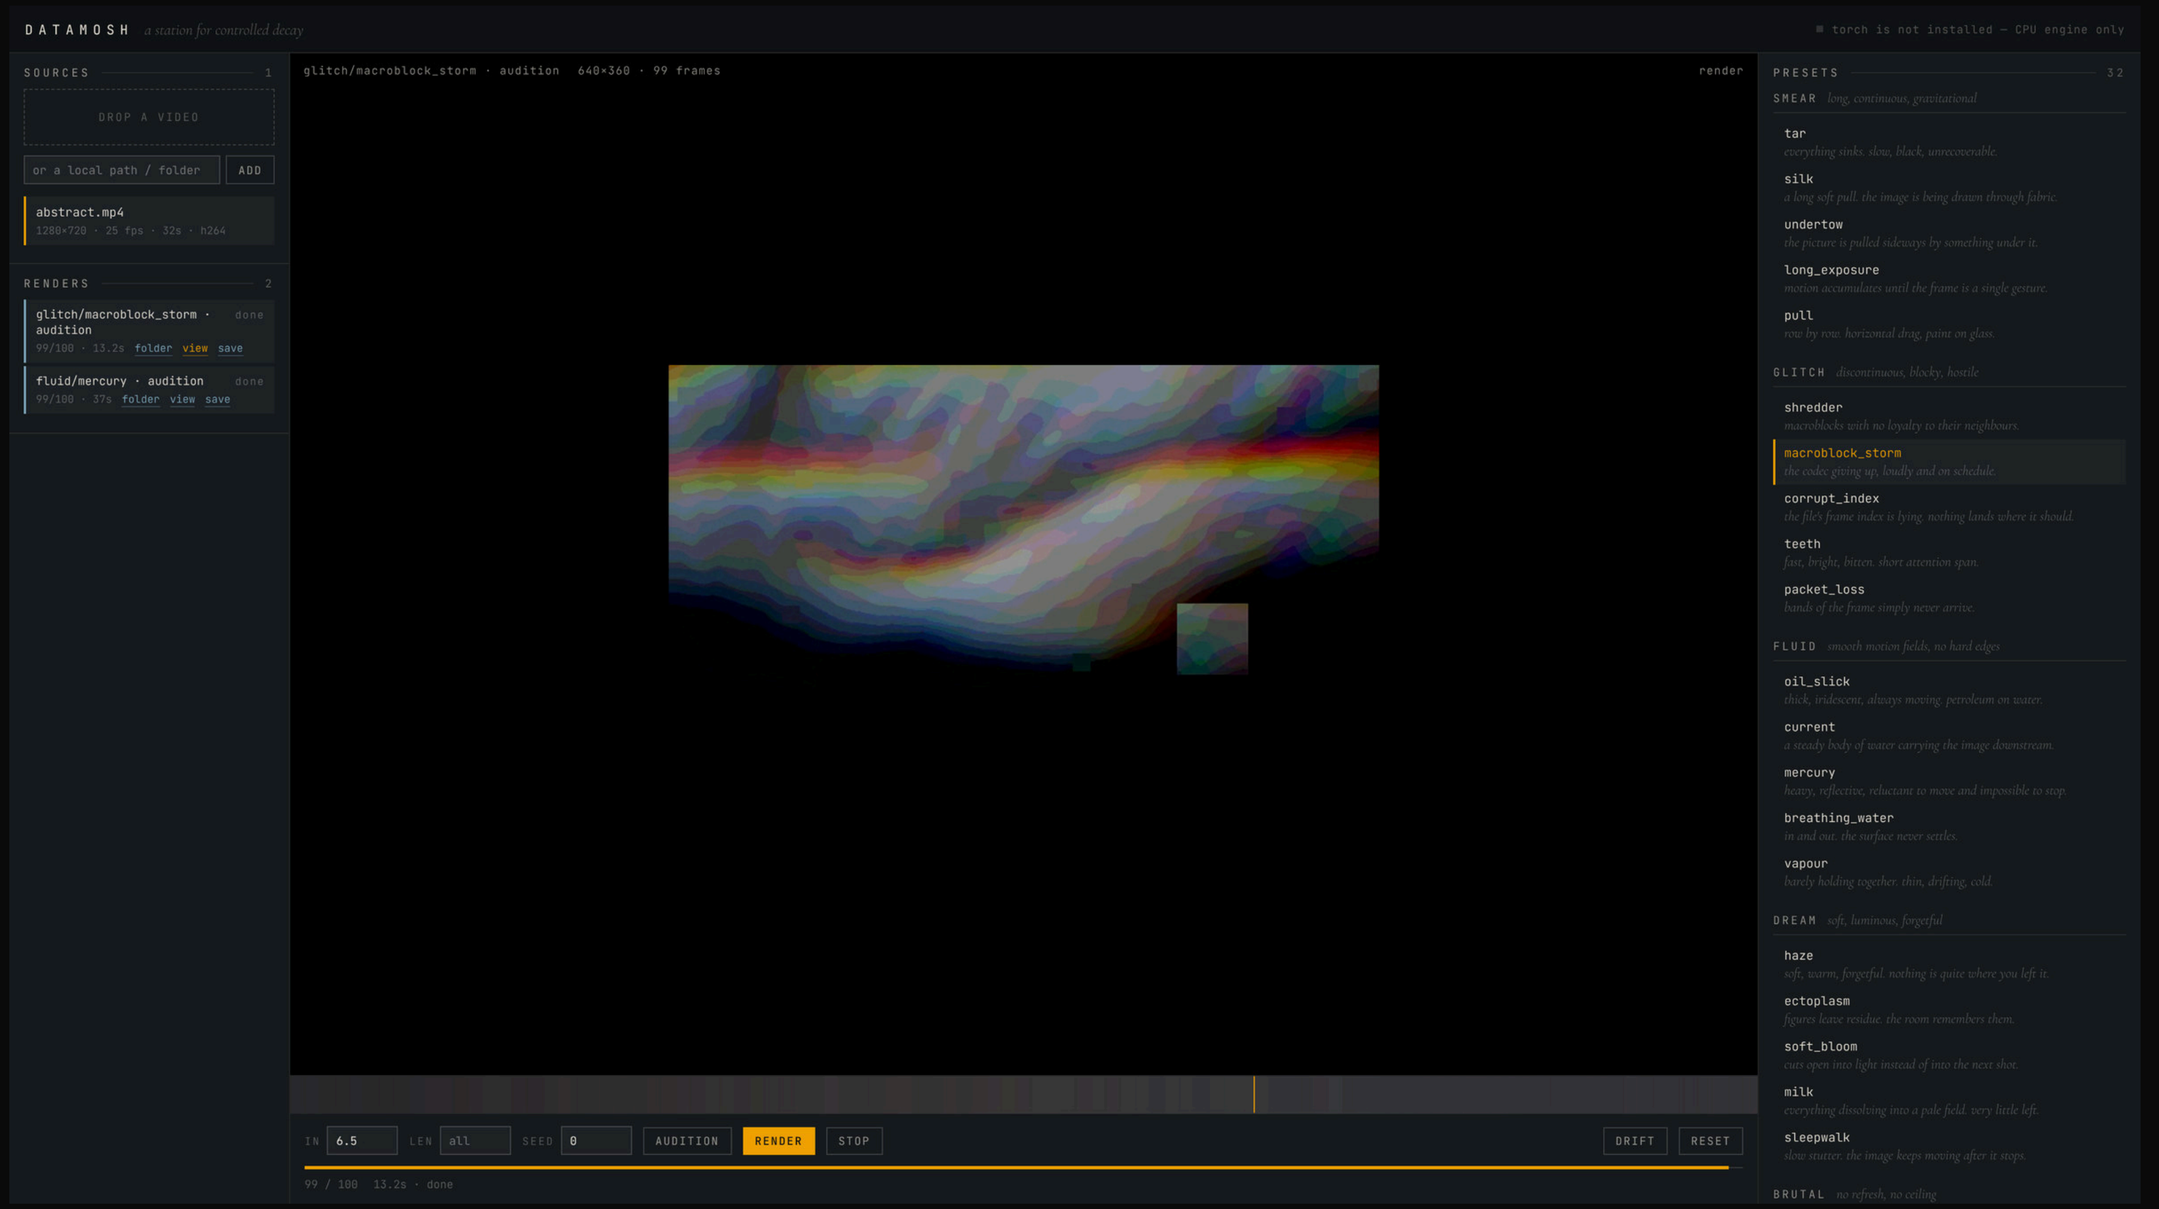
ectoplasm
figures leave residue. the room remembers them.

soft_bloom
cuts open into light instead of into the next shot.

milk
everything dissolving into a pale field. very little left.

sleepwalk
slow stutter. the image keeps moving after it stops.

BRUTAL *no refresh, no ceiling*


DATAMOSH
a station for controlled decay
torch is not installed — CPU engine only
SOURCES 1
DROP A VIDEO
or a local path / folder
ADD
abstract.mp4
1280×720 · 25 fps · 32s · h264
RENDERS 2
glitch/macroblock_storm · audition done
99/100 · 13.2s folder view save
fluid/mercury · audition done
99/100 · 37s folder view save
glitch/macroblock_storm · audition 640×360 · 99 frames
render
IN 6.5
LEN all
SEED 0
AUDITION
RENDER
STOP
DRIFT
RESET
99 / 100 13.2s · done
PRESETS 32
SMEAR long, continuous, gravitational
tar
everything sinks. slow, black, unrecoverable.
silk
a long soft pull. the image is being drawn through fabric.
undertow
the picture is pulled sideways by something under it.
long_exposure
motion accumulates until the frame is a single gesture.
pull
row by row. horizontal drag, paint on glass.
GLITCH discontinuous, blocky, hostile
shredder
macroblocks with no loyalty to their neighbours.
macroblock_storm
the codec giving up, loudly and on schedule.
corrupt_index
the file's frame index is lying. nothing lands where it should.
teeth
fast, bright, bitten. short attention span.
packet_loss
bands of the frame simply never arrive.
FLUID smooth motion fields, no hard edges
oil_slick
thick, iridescent, always moving. petroleum on water.
current
a steady body of water carrying the image downstream.
mercury
heavy, reflective, reluctant to move and impossible to stop.
breathing_water
in and out. the surface never settles.
vapour
barely holding together. thin, drifting, cold.
DREAM soft, luminous, forgetful
haze
soft, warm, forgetful. nothing is quite where you left it.
ectoplasm
figures leave residue. the room remembers them.
soft_bloom
cuts open into light instead of into the next shot.
milk
everything dissolving into a pale field. very little left.
sleepwalk
slow stutter. the image keeps moving after it stops.
BRUTAL no refresh, no ceiling

DATAMOSH *a station for controlled decay*

■ torch is not installed — CPU engine only

SOURCES 1

DROP A VIDEO

or a local path / folder | ADD

abstract.mp4
1280×720 · 25 fps · 32s · h264

RENDERS 3

dream/ectoplasm · audition done
100/100 · 21.7s folder view save

glitch/macroblock_storm · audition done
99/100 · 13.2s folder view save

fluid/mercury · audition done
99/100 · 37s folder view save

dream/ectoplasm · audition 640×360 · 100 frames

render



IN 16.2 | LEN all | SEED 0 | AUDITION | RENDER | STOP | DRIFT | RESET

100 / 100 21.7s · done

PRESETS 32

SMEAR *long, continuous, gravitational*

tar
everything sinks. slow, black, unrecoverable.

silk
a long soft pull. the image is being drawn through fabric.

undertow
the picture is pulled sideways by something under it.

long_exposure
motion accumulates until the frame is a single gesture.

pull
row by row. horizontal drag, paint on glass.

GLITCH *discontinuous, blocky, hostile*

shredder
macroblocks with no loyalty to their neighbours.

macroblock_storm
the codec giving up, loudly and on schedule.

corrupt_index
the file's frame index is lying. nothing lands where it should.

teeth
fast, bright, bitten. short attention span.

packet_loss
bands of the frame simply never arrive.

FLUID *smooth motion fields, no hard edges*

oil_slick
thick, iridescent, always moving. petroleum on water.

current
a steady body of water carrying the image downstream.

mercury
heavy, reflective, reluctant to move and impossible to stop.

breathing_water
in and out. the surface never settles.

vapour
barely holding together. thin, drifting, cold.

DREAM *soft, luminous, forgetful*

haze
soft, warm, forgetful. nothing is quite where you left it.

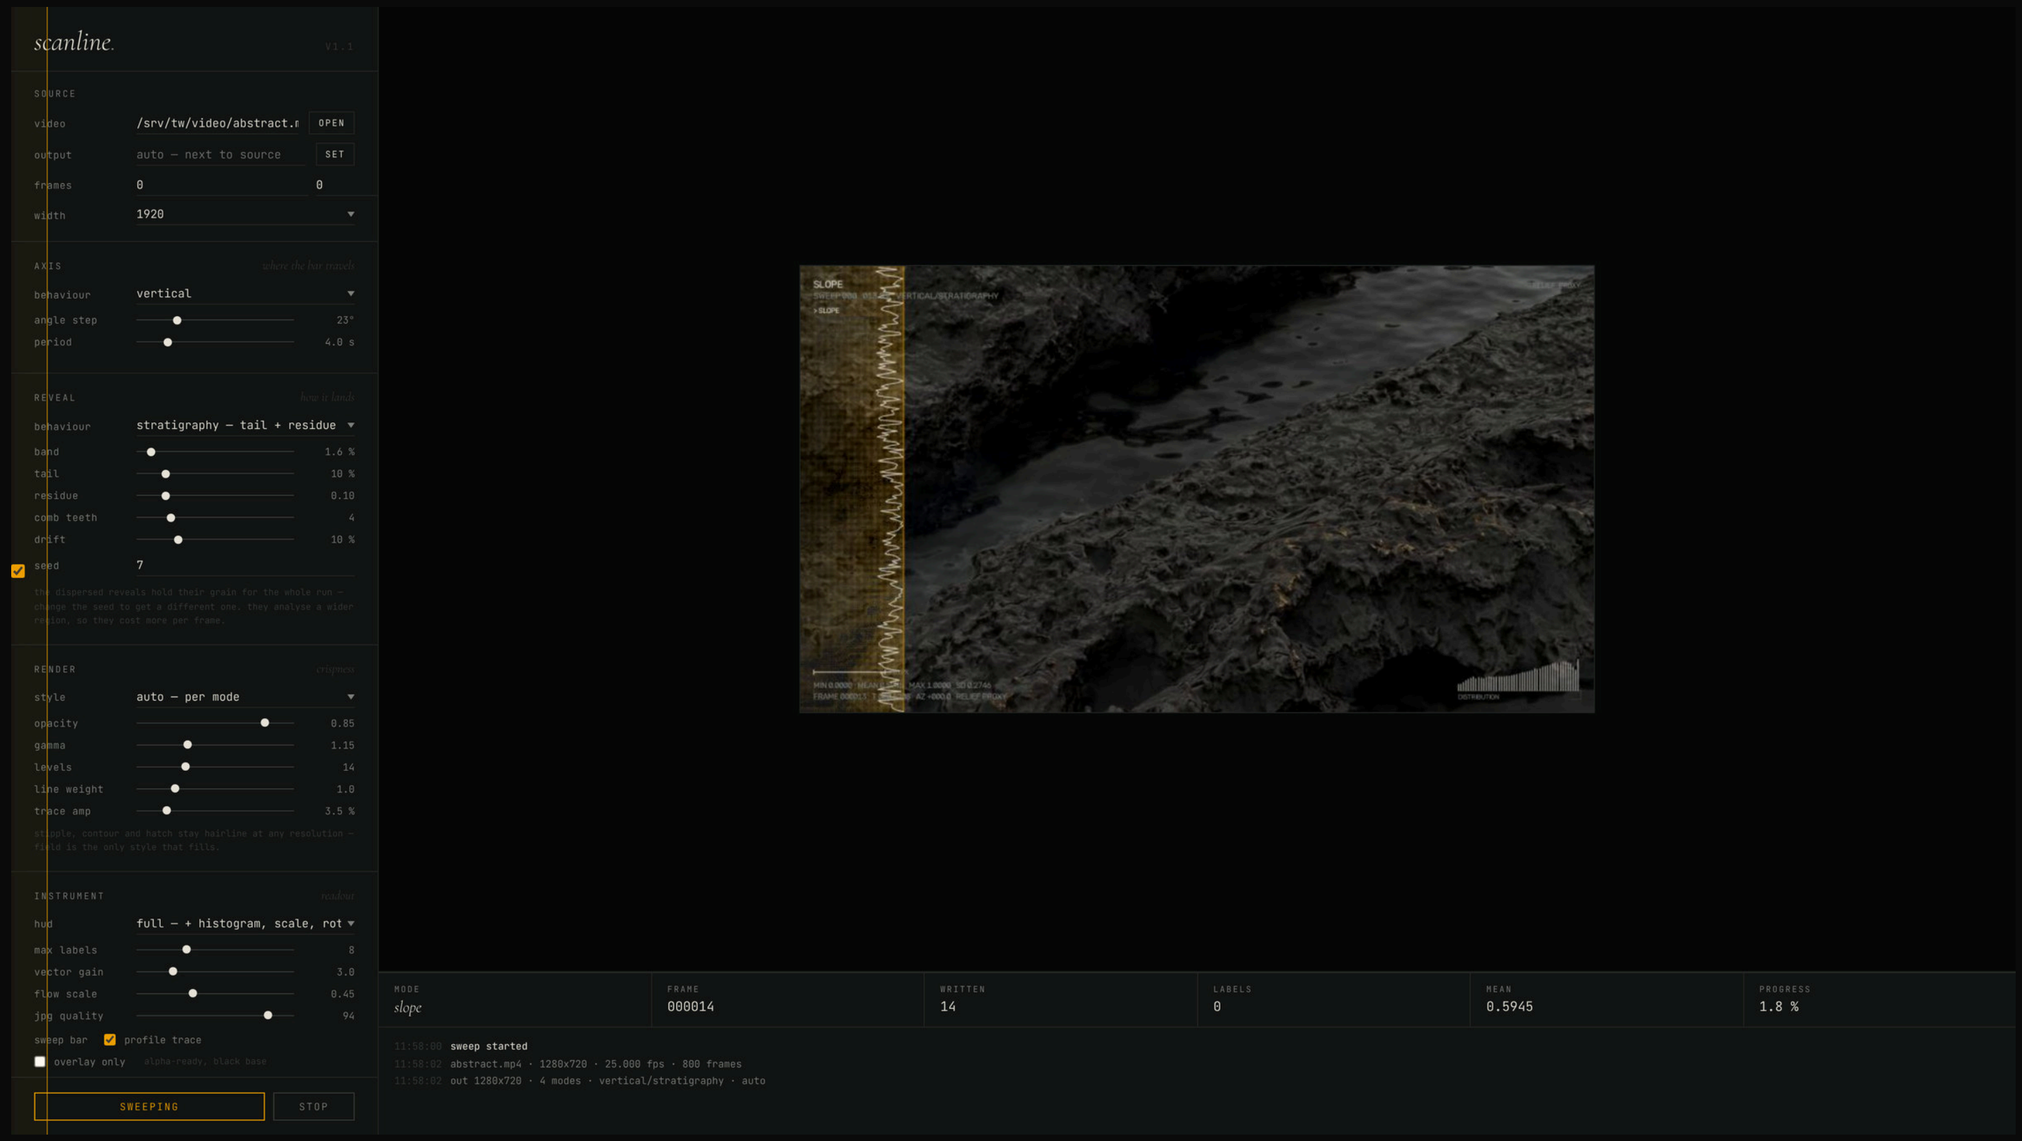
ectoplasm
figures leave residue. the room remembers them.

soft_bloom
cuts open into light instead of into the next shot.

milk
everything dissolving into a pale field. very little left.

sleepwalk
slow stutter. the image keeps moving after it stops.

BRUTAL *no refresh, no ceiling*



scanline

2026

MOVING IMAGE

01 · texture variance

T09

T09

scanline

2026

FIELD	Moving image — process, analysis, capture
VERSION	1.1
TECHNOLOGY	Python (numpy, OpenCV, optional depth model), browser UI, ffmpeg
SOURCE	scanline_v1.1.zip

A sweeping bar crosses the moving image and reveals, behind it, one of 35 analytical readings — terrain (slope, curvature, ruggedness), photometric (glint, water index, haze), motion (flow, divergence, vorticity, time-to-contact), detection (landmass, horizon, rectilinear anomalies) and depth. The frame is thus never fully in one epistemology. Output: an image sequence, a preview film and a per-frame scan log.

scanline. V1.1

SOURCE

video /srv/tw/video/abstract.m OPEN

output auto — next to source SET

frames 0 0

width 1920

AXIS — where the bar travels

behaviour vertical

angle step 23°

period 4.0 s

REVEAL — how it lands

behaviour stratigraphy — tail + residue

band 1.6 %

tail 10 %

residue 0.10

comb teeth 4

drift 10 %

seed 7

the dispersed reveals hold their grain for the whole run — change the seed to get a different one. they analyse a wider region, so they cost more per frame.

RENDER — crispness

style auto — per mode

opacity 0.85

gamma 1.15

levels 14

line weight 1.0

trace amp 3.5 %

stipple, contour and hatch stay hairline at any resolution — field is the only style that fills.

INSTRUMENT — readout

hud full — + histogram, scale, rot

max labels 8

vector gain 3.0

flow scale 0.45

jpg quality 94

sweep bar profile trace

overlay only alpha-ready, black base

SWEEPING STOP



MODE	FRAME	WRITTEN	LABELS	MEAN	PROGRESS
slope	000004	4	0	0.5896	0.5 %

11:58:00 sweep started
11:58:02 abstract.mp4 · 1280x720 · 25.000 fps · 800 frames
11:58:02 out 1280x720 · 4 modes · vertical/stratigraphy · auto

scanline. V1.1

SOURCE

video /srv/tw/video/abstract.m OPEN
output auto — next to source SET
frames 0 0
width 1920

AXIS *where the bar travels*

behaviour vertical
angle step 23°
period 4.0 s

REVEAL *how it lands*

behaviour stratigraphy — tail + residue
band 1.6 %
tail 10 %
residue 0.10
comb teeth 4
drift 10 %
seed 7

the dispersed reveals hold their grain for the whole run — change the seed to get a different one. they analyse a wider region, so they cost more per frame.

RENDER *crispness*

style auto — per mode
opacity 0.85
gamma 1.15
levels 14
line weight 1.0
trace amp 3.5 %

stipple, contour and hatch stay hairline at any resolution — field is the only style that fills.

INSTRUMENT *readout*

hud full — + histogram, scale, rot
max labels 8
vector gain 3.0
flow scale 0.45
jpg quality 94
sweep bar profile trace
overlay only *alpha-ready, black base*

SWEEPING STOP



MODE	FRAME	WRITTEN	LABELS	MEAN	PROGRESS
slope	000024	24	0	0.5916	3.0 %

11:58:00 **sweep started**
11:58:02 abstract.mp4 · 1280x720 · 25.000 fps · 800 frames
11:58:02 out 1280x720 · 4 modes · vertical/stratigraphy · auto

scanline. V1.1

SOURCE

video /srv/tw/video/abstract.m OPEN

output auto — next to source SET

frames 0 0

width 1920

AXIS — where the bar travels

behaviour vertical

angle step 23°

period 4.0 s

REVEAL — how it lands

behaviour stratigraphy — tail + residue

band 1.6 %

tail 10 %

residue 0.10

comb teeth 4

drift 10 %

seed 7

the dispersed reveals hold their grain for the whole run — change the seed to get a different one. they analyse a wider region, so they cost more per frame.

RENDER — crispness

style auto — per mode

opacity 0.85

gamma 1.15

levels 14

line weight 1.0

trace amp 3.5 %

stipple, contour and hatch stay hairline at any resolution — field is the only style that fills.

INSTRUMENT — readout

hud full — + histogram, scale, rot

max labels 8

vector gain 3.0

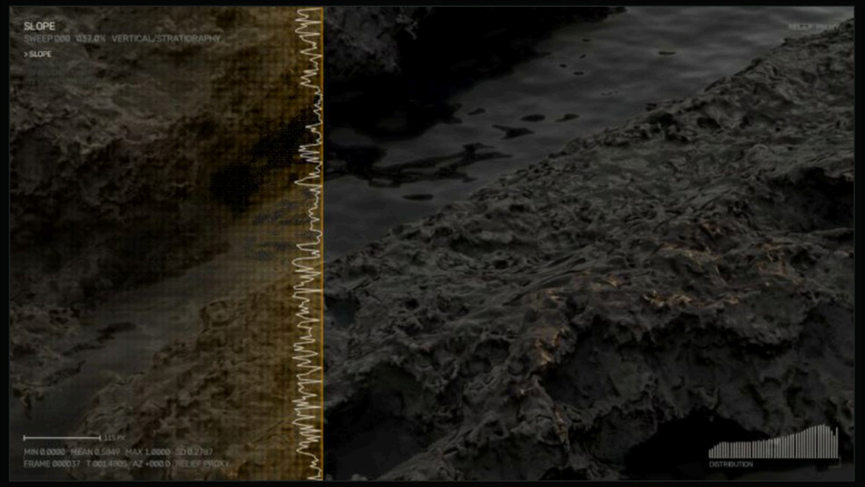
flow scale 0.45

jpg quality 94

sweep bar profile trace

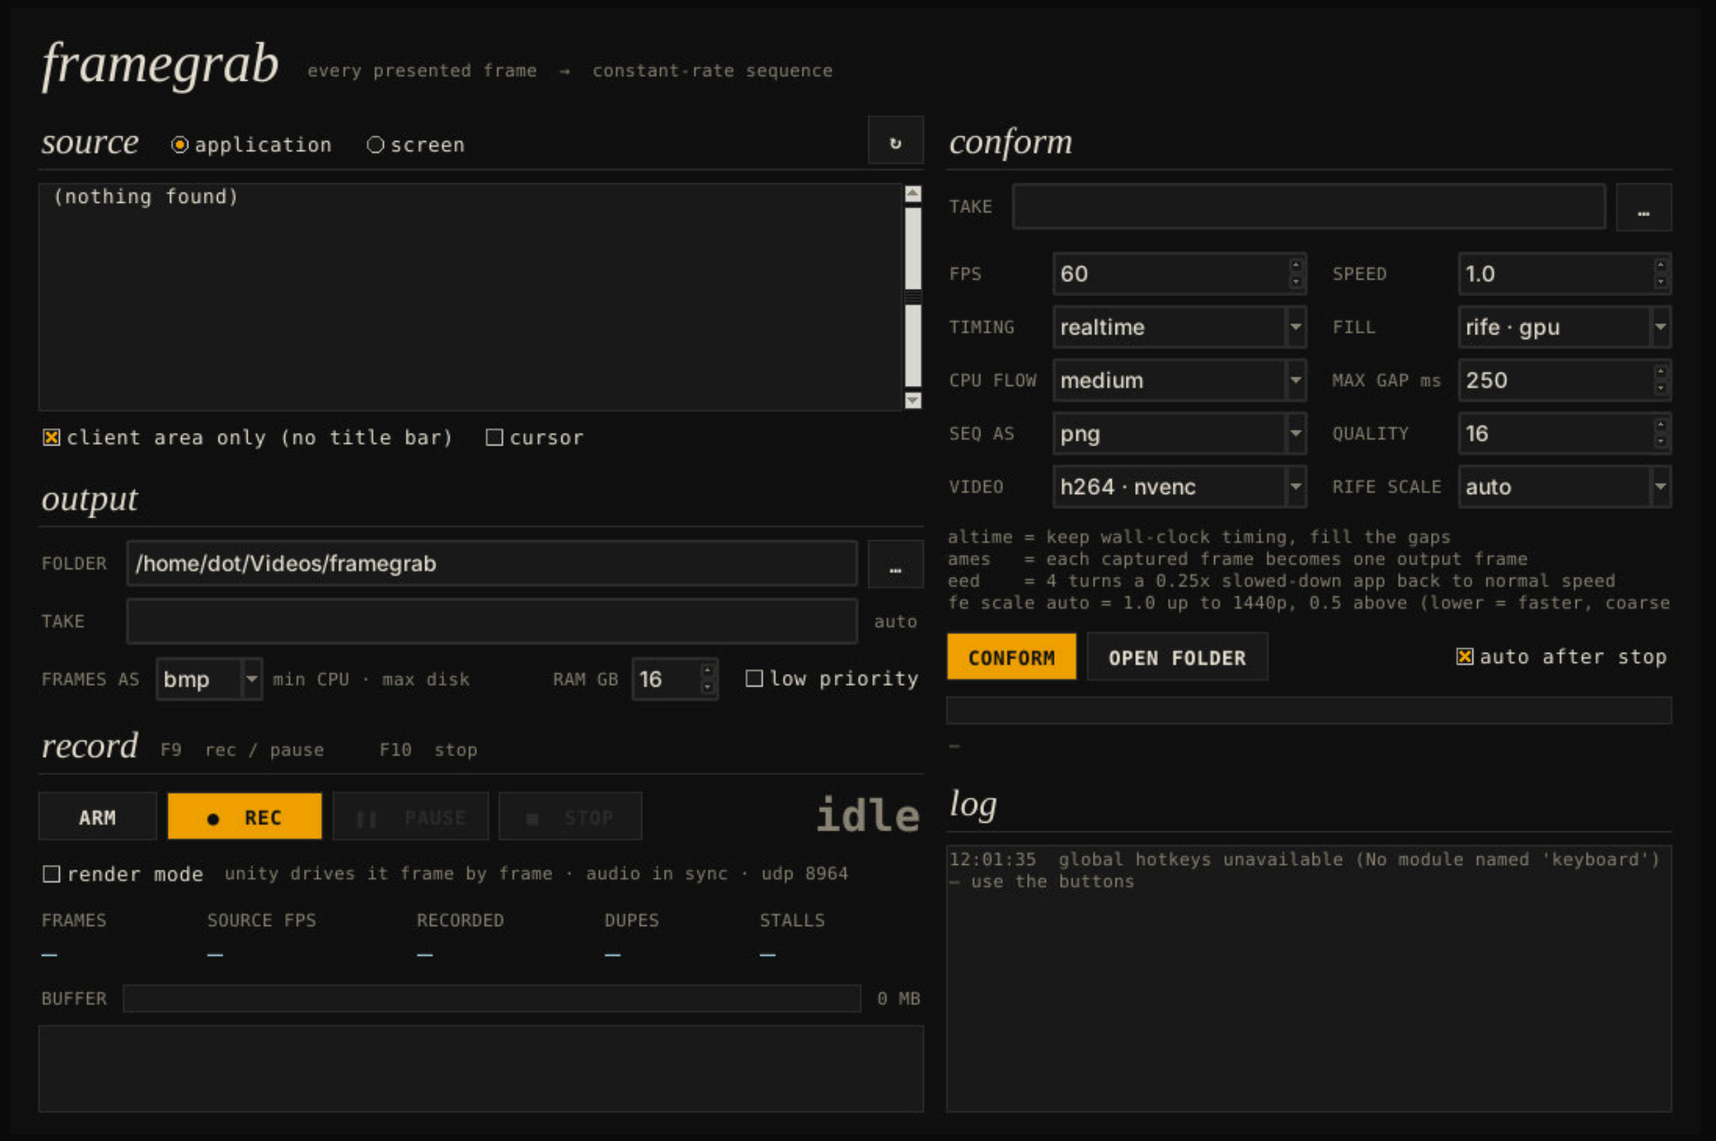
overlay only alpha-ready, black base

SWEEPING STOP



MODE	FRAME	WRITTEN	LABELS	MEAN	PROGRESS
slope	000038	38	0	0.5849	4.8 %

11:58:00 sweep started
11:58:02 abstract.mp4 · 1280x720 · 25.000 fps · 800 frames
11:58:02 out 1280x720 · 4 modes · vertical/stratigraphy · auto


framegrab
every presented frame → constant-rate sequence
source
application
screen
(nothing found)
client area only (no title bar)
cursor
output
FOLDER
/home/dot/Videos/framegrab
TAKE
auto
FRAMES AS
bmp
min CPU · max disk
RAM GB
16
low priority
record
F9 rec / pause
F10 stop
ARM
REC
PAUSE
STOP
idle
render mode
unity drives it frame by frame · audio in sync · udp 8964
FRAMES
SOURCE FPS
RECORDED
DUPES
STALLS
BUFFER
0 MB
conform
TAKE
FPS
60
SPEED
1.0
TIMING
realtime
FILL
rife · gpu
CPU FLOW
medium
MAX GAP ms
250
SEQ AS
png
QUALITY
16
VIDEO
h264 · nvenc
RIFE SCALE
auto
altime = keep wall-clock timing, fill the gaps
ames = each captured frame becomes one output frame
eed = 4 turns a 0.25x slowed-down app back to normal speed
fe scale auto = 1.0 up to 1440p, 0.5 above (lower = faster, coarse
CONFORM
OPEN FOLDER
auto after stop
log
12:01:35 global hotkeys unavailable (No module named 'keyboard')
– use the buttons

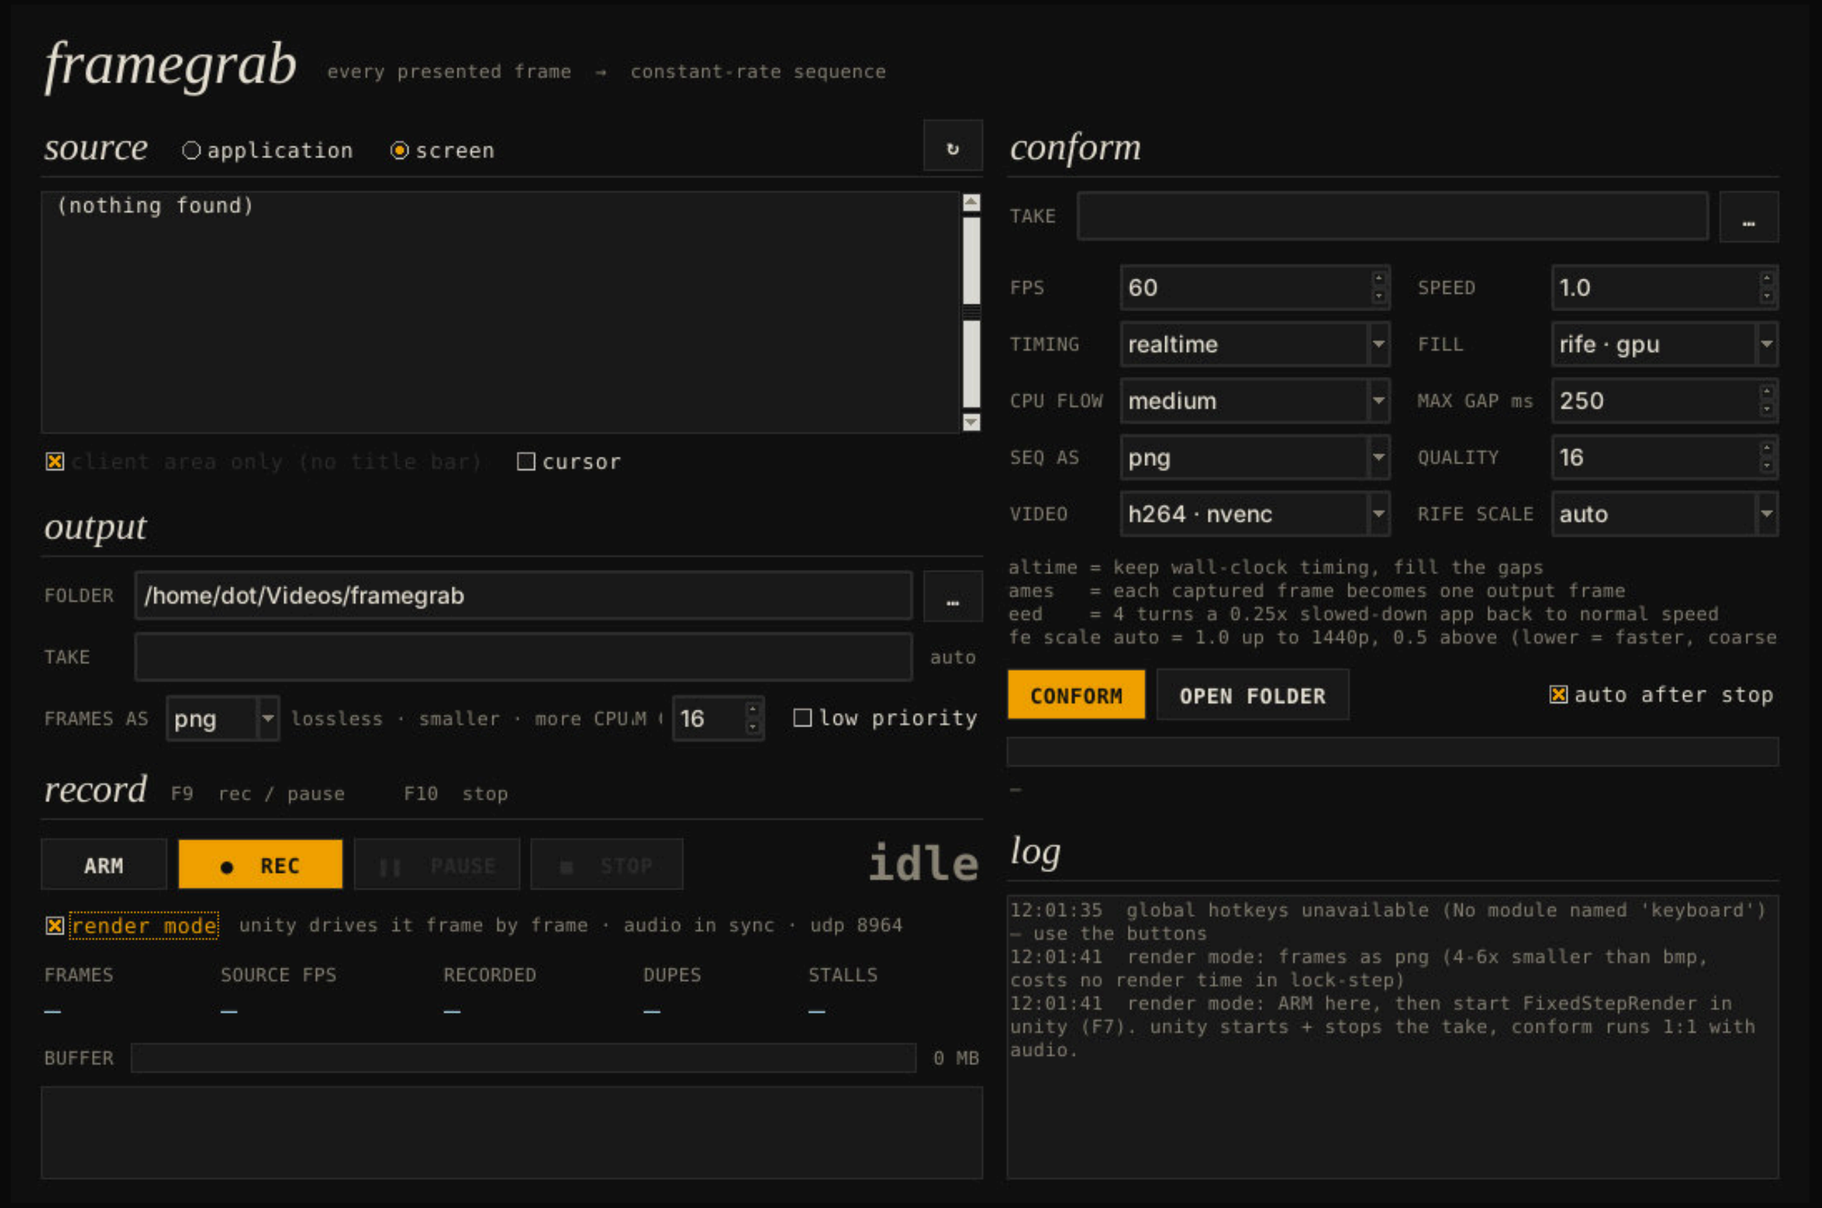
T10

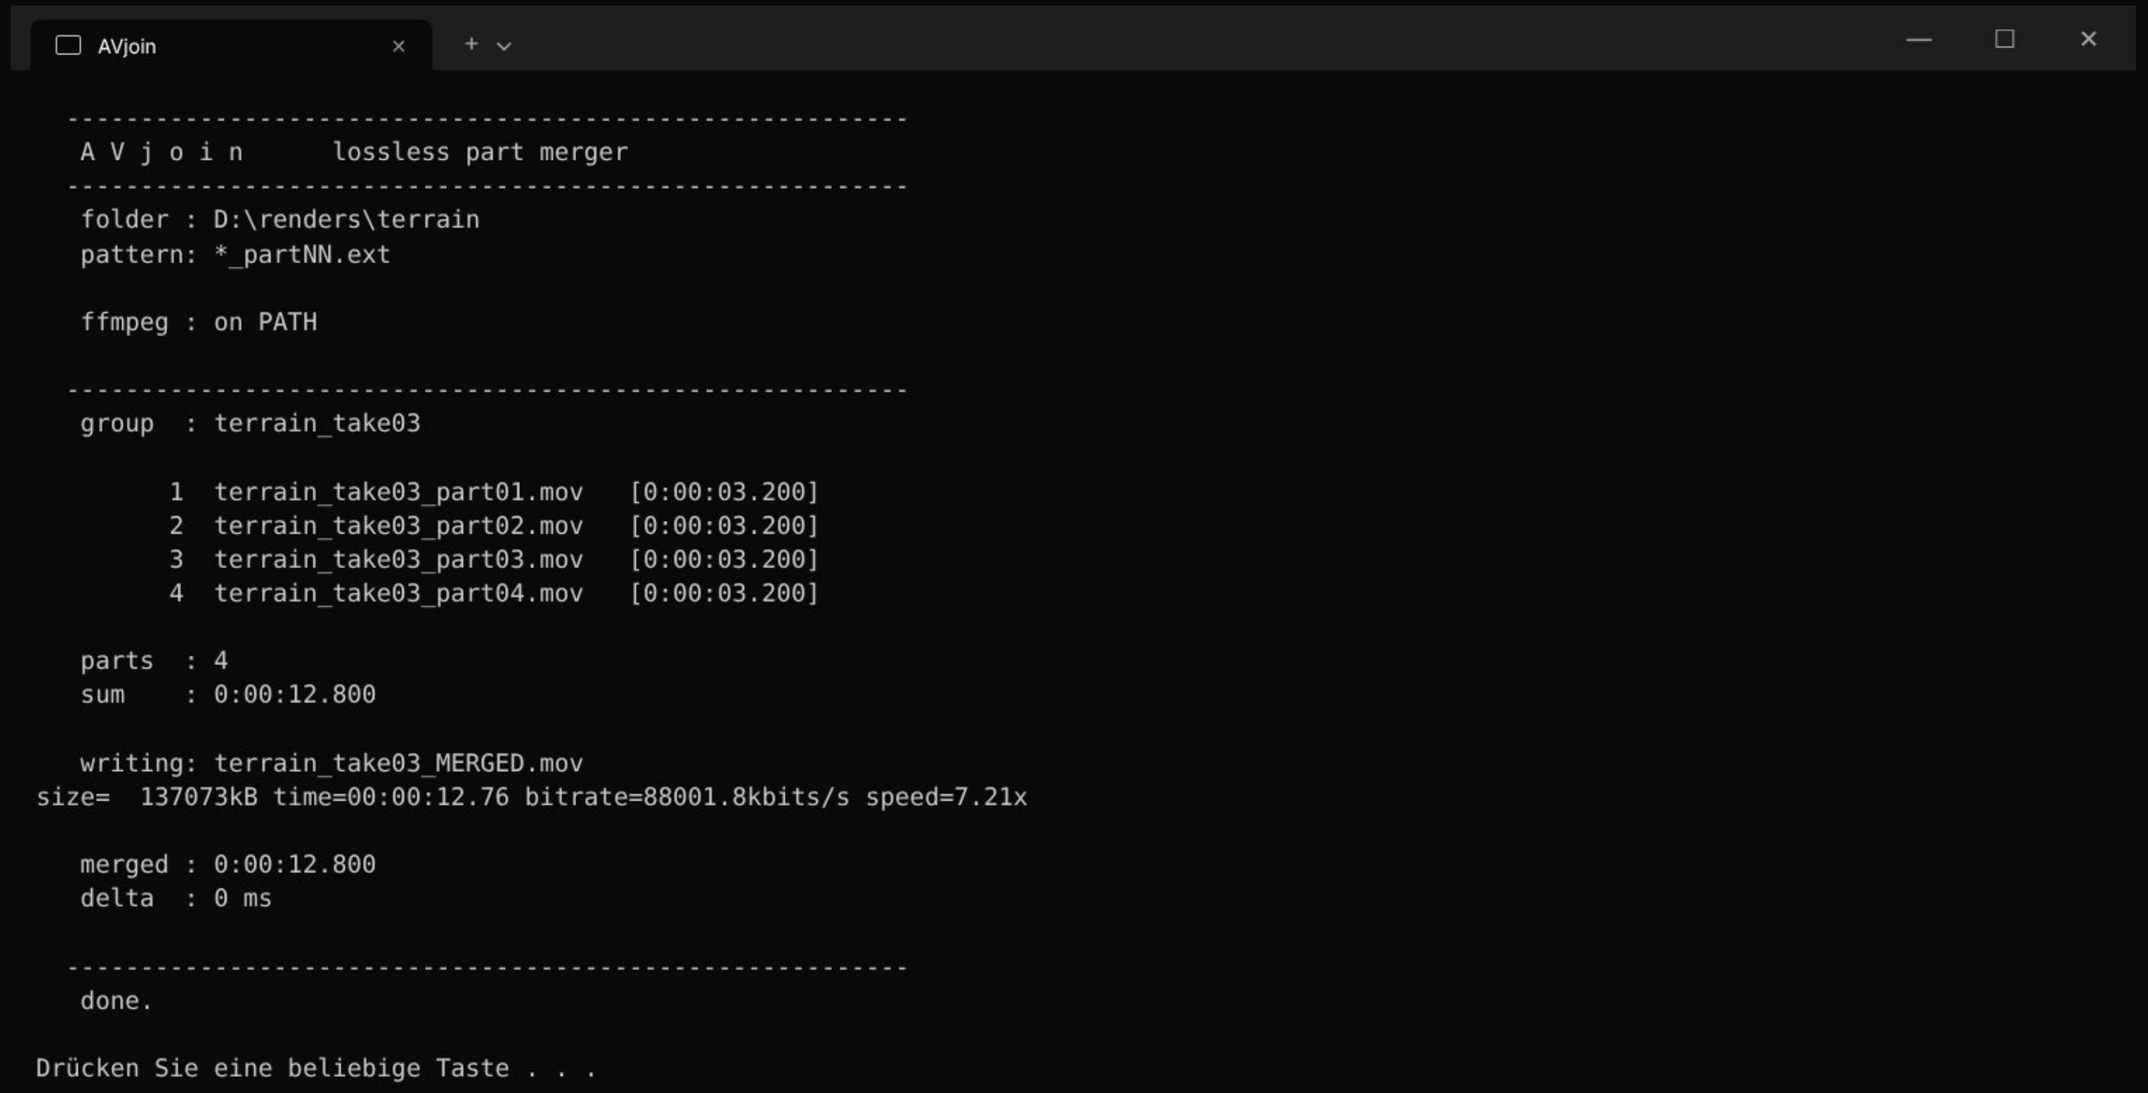
framegrab

2026

FIELD	Moving image — process, analysis, capture
VERSION	fixed-step
TECHNOLOGY	Python/tkinter, Windows Graphics Capture, RIFE (GPU) interpolation, ffmpeg; Unity fixed-step render bridge (C#)
SOURCE	260918_FrameGrabberFixedStep_noIntAudio.zip

Captures every frame an application actually presents — including irregular, frame-generated or lagging sources — with high-resolution timestamps, and conforms them to an even timeline (hold, blend, optical flow or RIFE interpolation). A render mode lets Unity drive capture frame by frame with synchronised audio over UDP, turning real-time engines into deterministic offline renderers.


framegrab
every presented frame → constant-rate sequence
source
application
screen
(nothing found)
client area only (no title bar)
cursor
output
FOLDER
/home/dot/Videos/framegrab
TAKE
auto
FRAMES AS
png
lossless · smaller · more CPU
16
low priority
record
F9 rec / pause
F10 stop
ARM
REC
PAUSE
STOP
idle
render mode
unity drives it frame by frame · audio in sync · udp 8964
FRAMES
SOURCE FPS
RECORDED
DUPES
STALLS
BUFFER
0 MB
conform
TAKE
FPS
60
SPEED
1.0
TIMING
realtime
FILL
rife · gpu
CPU FLOW
medium
MAX GAP ms
250
SEQ AS
png
QUALITY
16
VIDEO
h264 · nvenc
RIFE SCALE
auto
altime = keep wall-clock timing, fill the gaps
ames = each captured frame becomes one output frame
eed = 4 turns a 0.25x slowed-down app back to normal speed
fe scale auto = 1.0 up to 1440p, 0.5 above (lower = faster, coarse
CONFORM
OPEN FOLDER
auto after stop
log
12:01:35 global hotkeys unavailable (No module named 'keyboard') — use the buttons
12:01:41 render mode: frames as png (4-6x smaller than bmp, costs no render time in lock-step)
12:01:41 render mode: ARM here, then start FixedStepRender in unity (F7). unity starts + stops the take, conform runs 1:1 with audio.



AVjoin

2026

MOVING IMAGE

01 · merge run (reconstructed on test parts)

T11

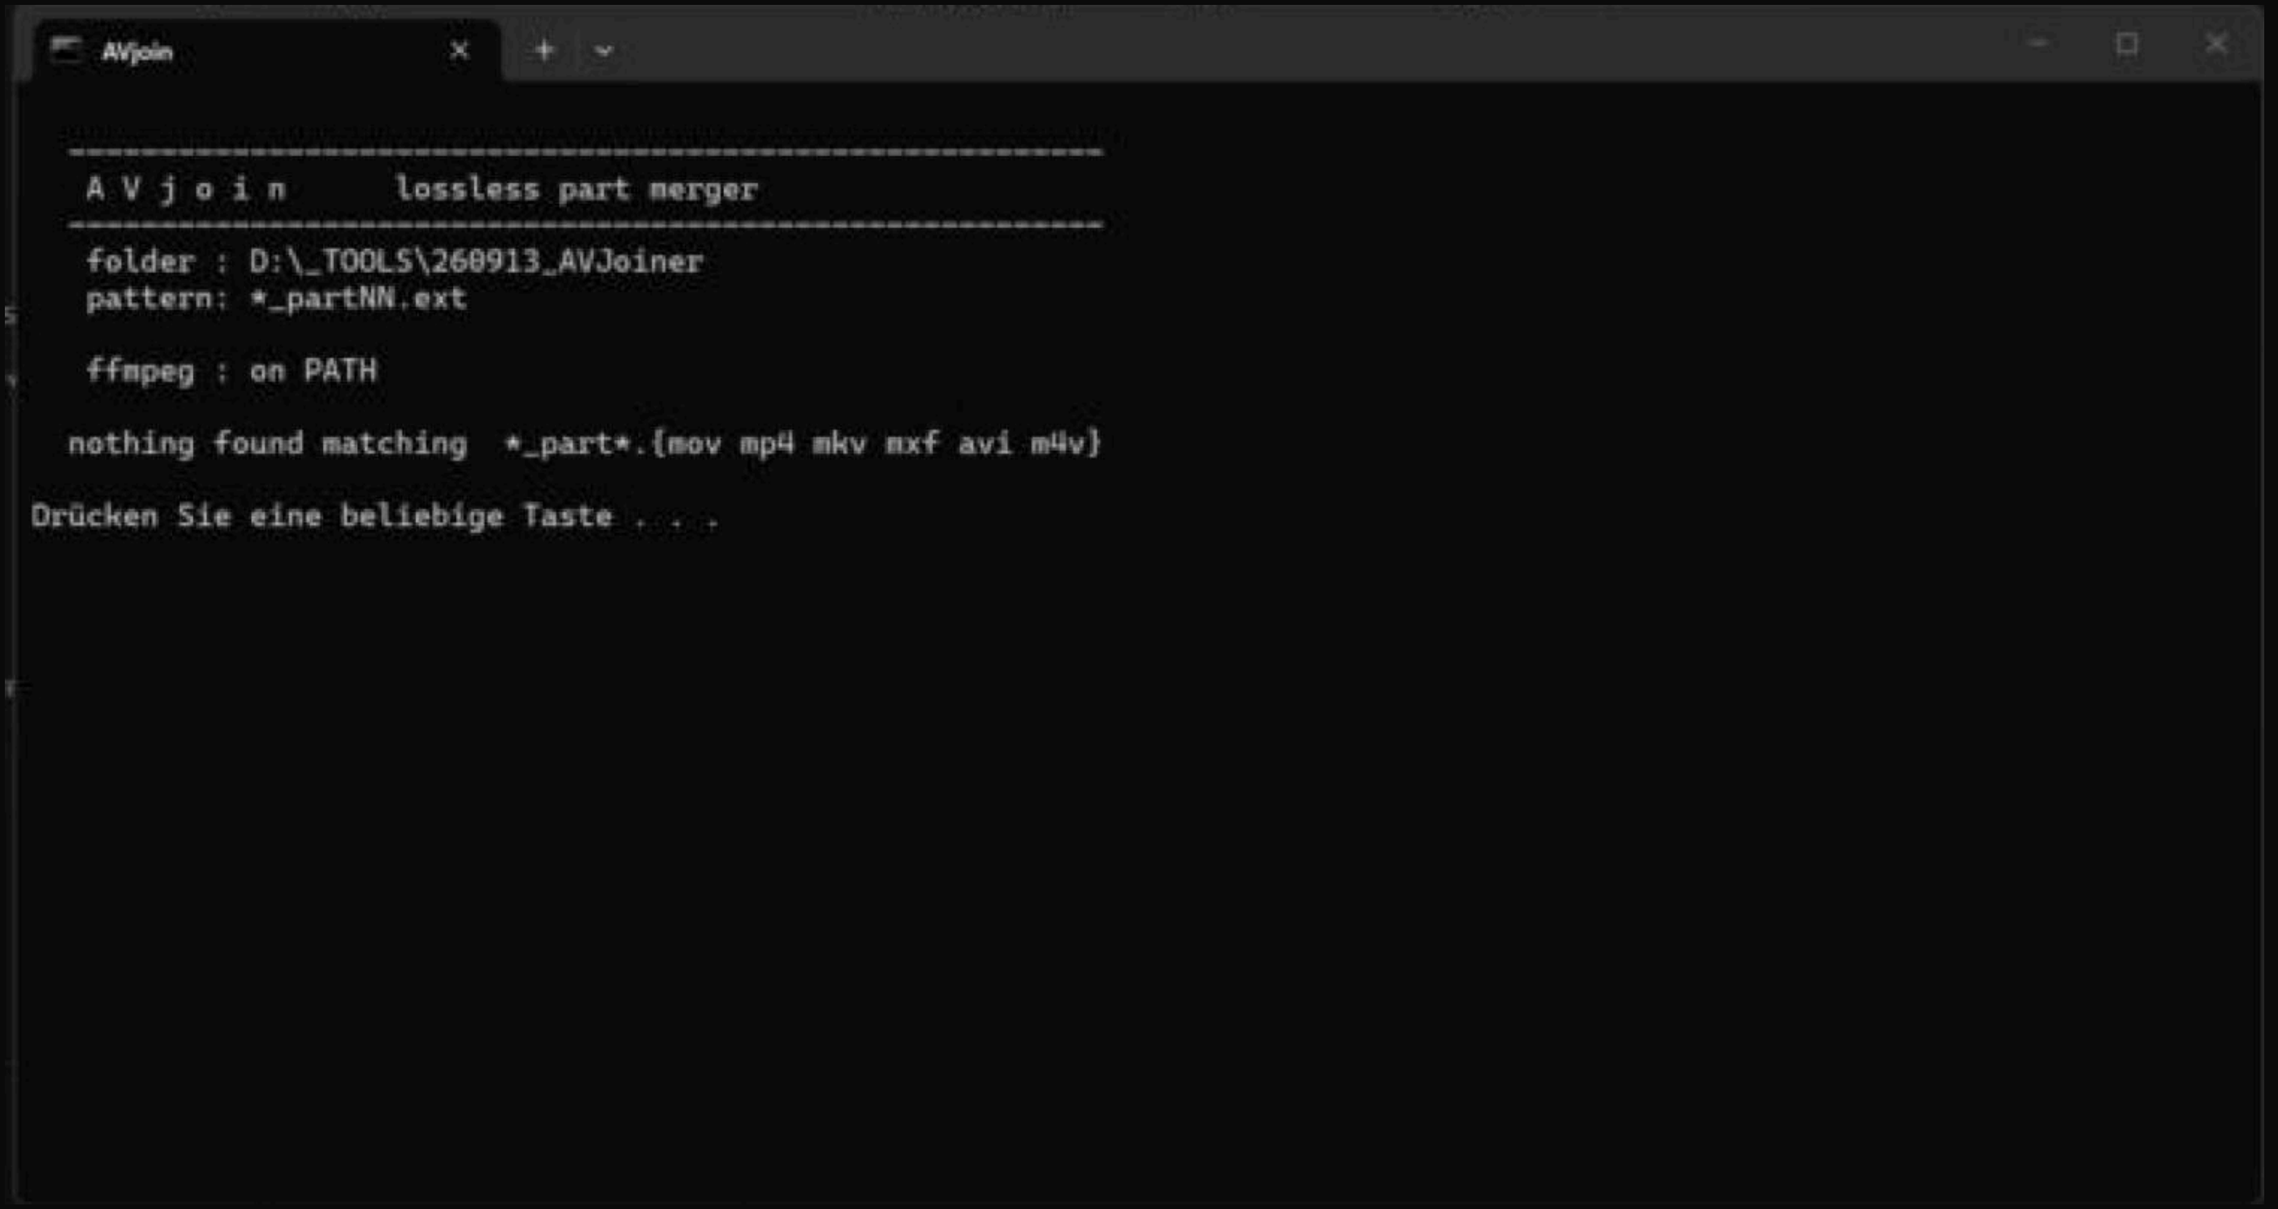
T11

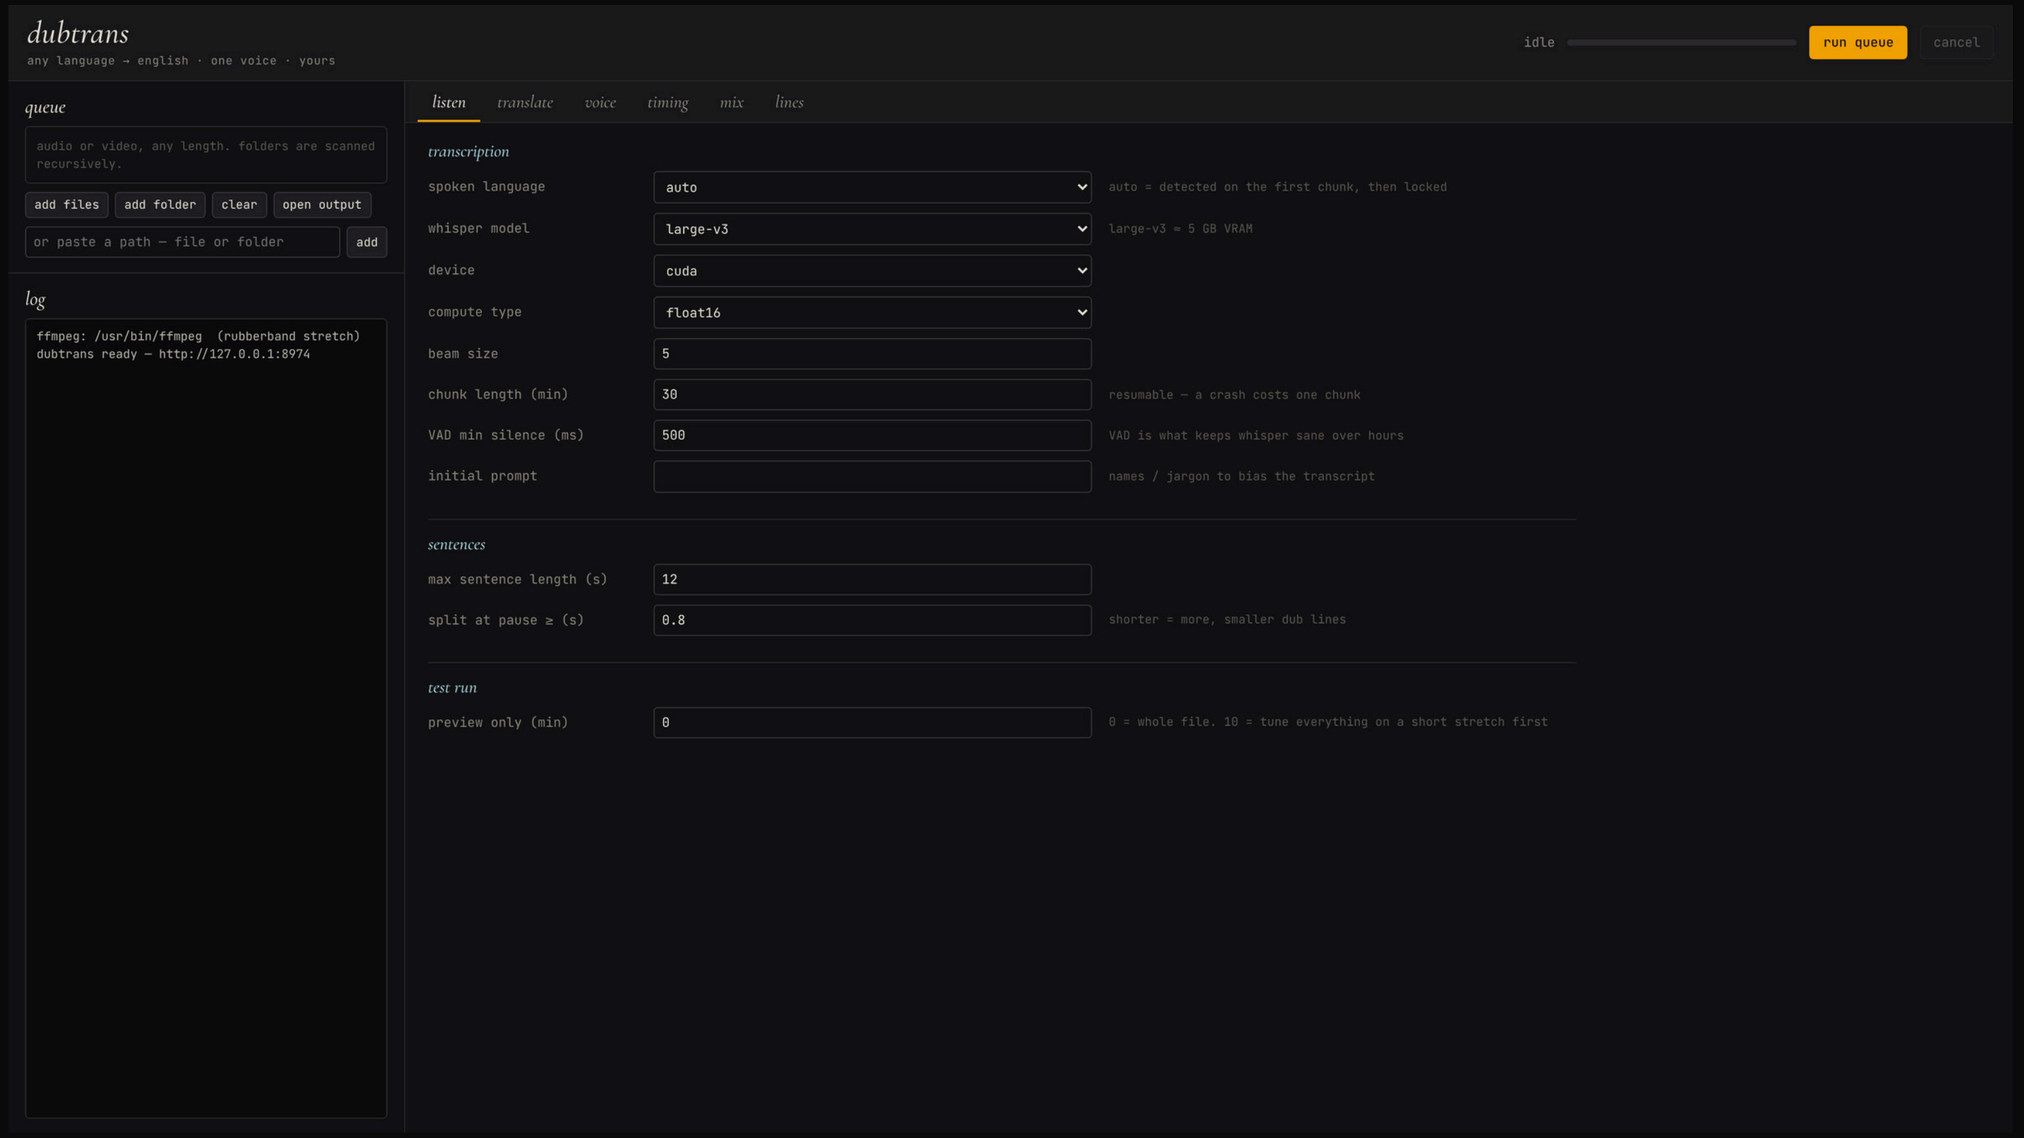
AVjoin

2026

FIELD	Moving image — process, analysis, capture
TECHNOLOGY	Windows batch, ffmpeg/ffprobe
SOURCE	260913_AVjoin.bat

Lossless concatenation of numbered video parts (name_partNN) by stream copy: groups by base name, sorts numerically, verifies identical stream parameters across parts and checks the merged duration against the sum of the parts. Counterpart to AVsplitter.


AVjoin
A V j o i n lossless part merger
folder : D:_TOOLS\260913_AVJoiner
pattern: *_partNN.ext
ffmpeg : on PATH
nothing found matching *_part*.{mov mp4 mkv mxf avi m4v}
Drücken Sie eine beliebige Taste . . .


dubtrans
any language → english · one voice · yours
idle
run queue
cancel
queue
audio or video, any length. folders are scanned recursively.
add files
add folder
clear
open output
or paste a path — file or folder
add
log
ffmpeg: /usr/bin/ffmpeg (rubberband stretch)
dubtrans ready — http://127.0.0.1:8974
listen
translate
voice
timing
mix
lines
transcription
spoken language
auto
auto = detected on the first chunk, then locked
whisper model
large-v3
large-v3 ≈ 5 GB VRAM
device
cuda
compute type
float16
beam size
5
chunk length (min)
30
resumable — a crash costs one chunk
VAD min silence (ms)
500
VAD is what keeps whisper sane over hours
initial prompt
names / jargon to bias the transcript
sentences
max sentence length (s)
12
split at pause ≥ (s)
0.8
shorter = more, smaller dub lines
test run
preview only (min)
0
0 = whole file. 10 = tune everything on a short stretch first

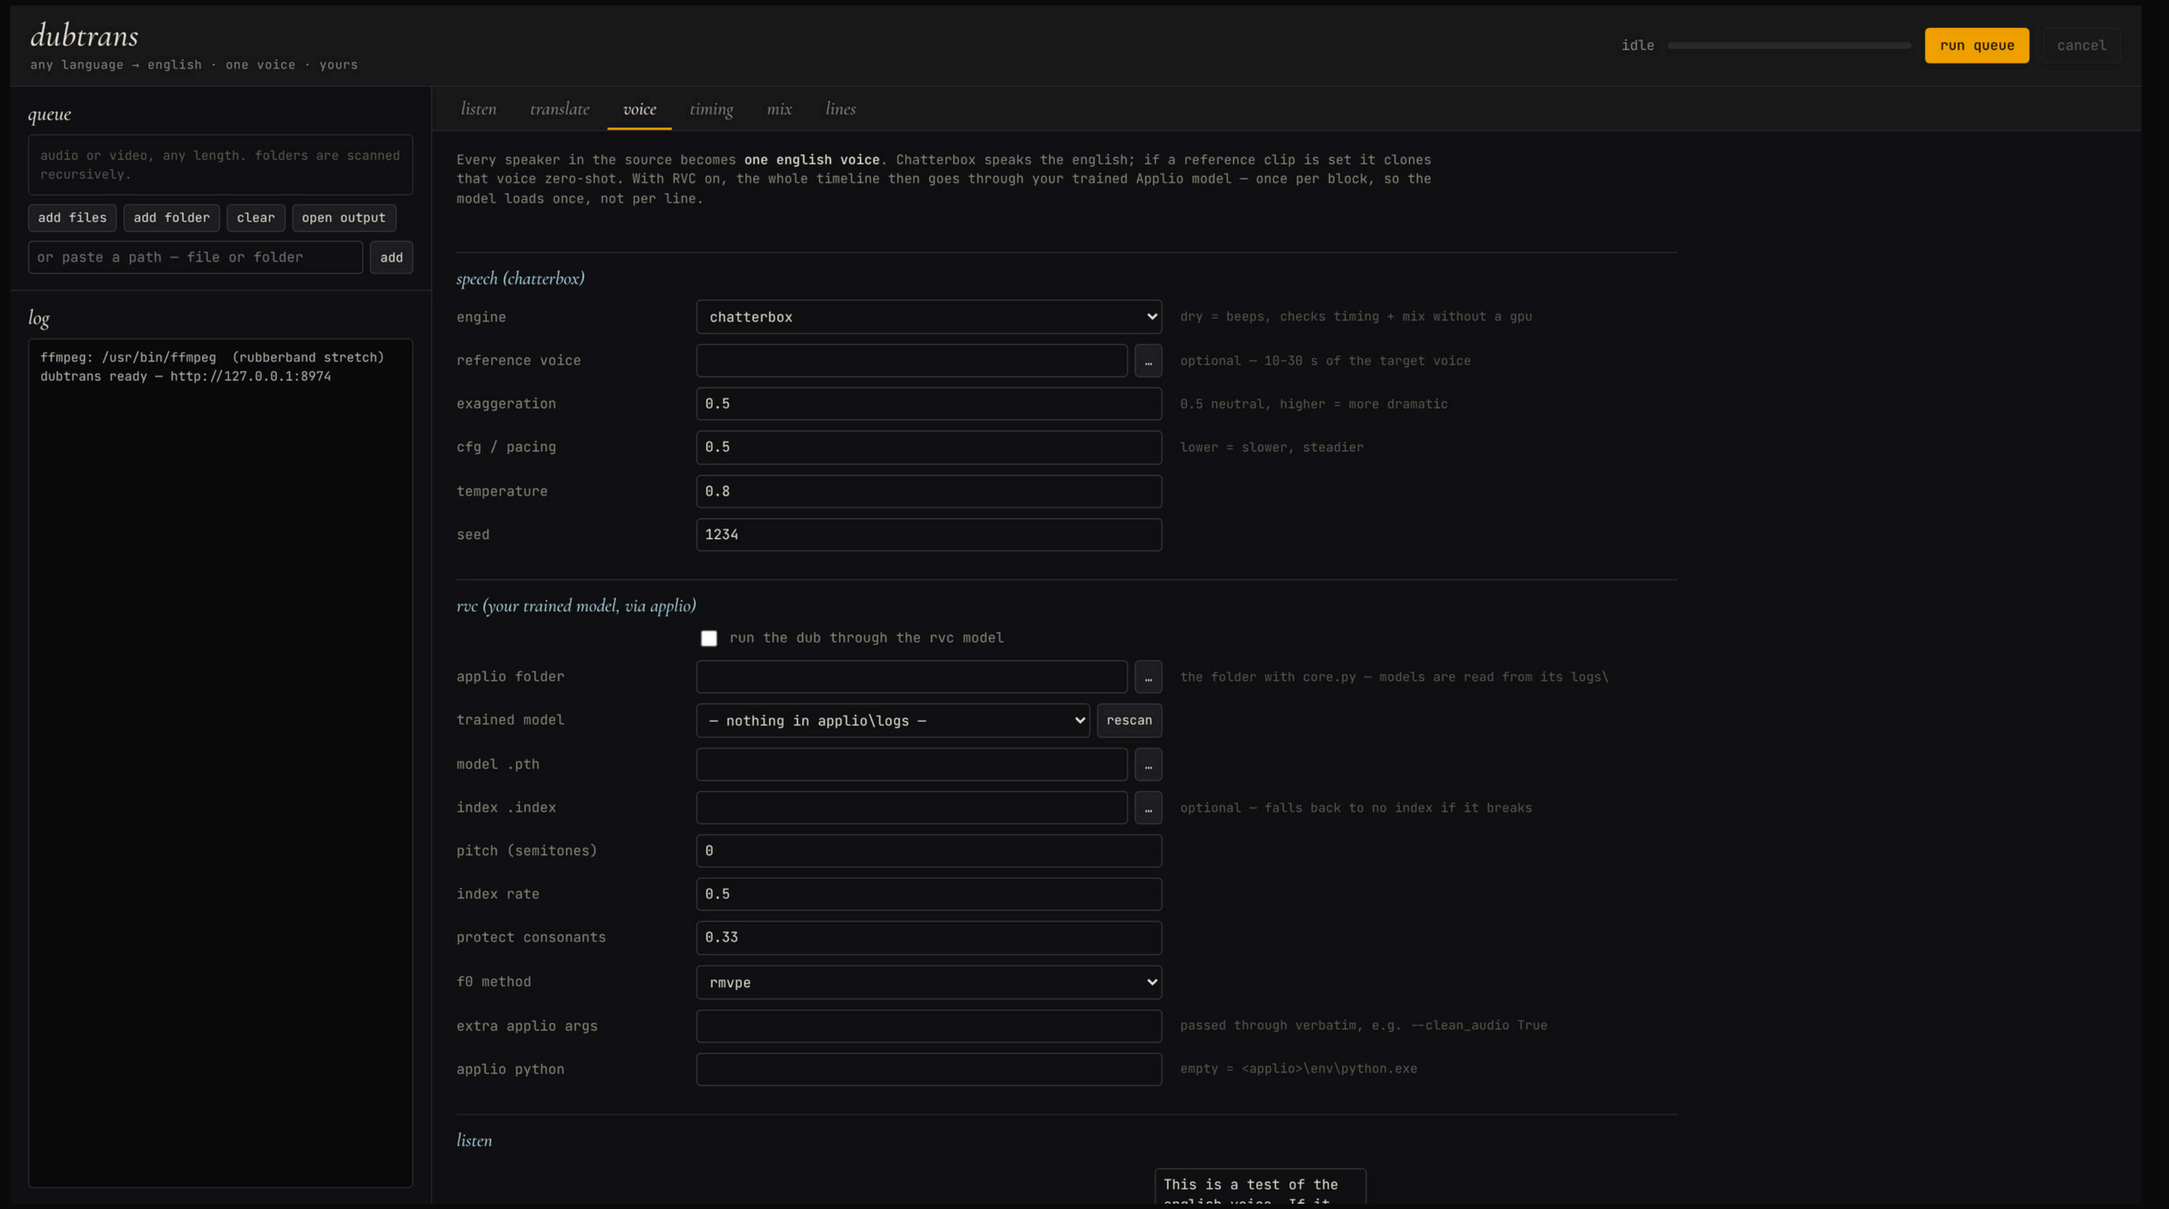
T12

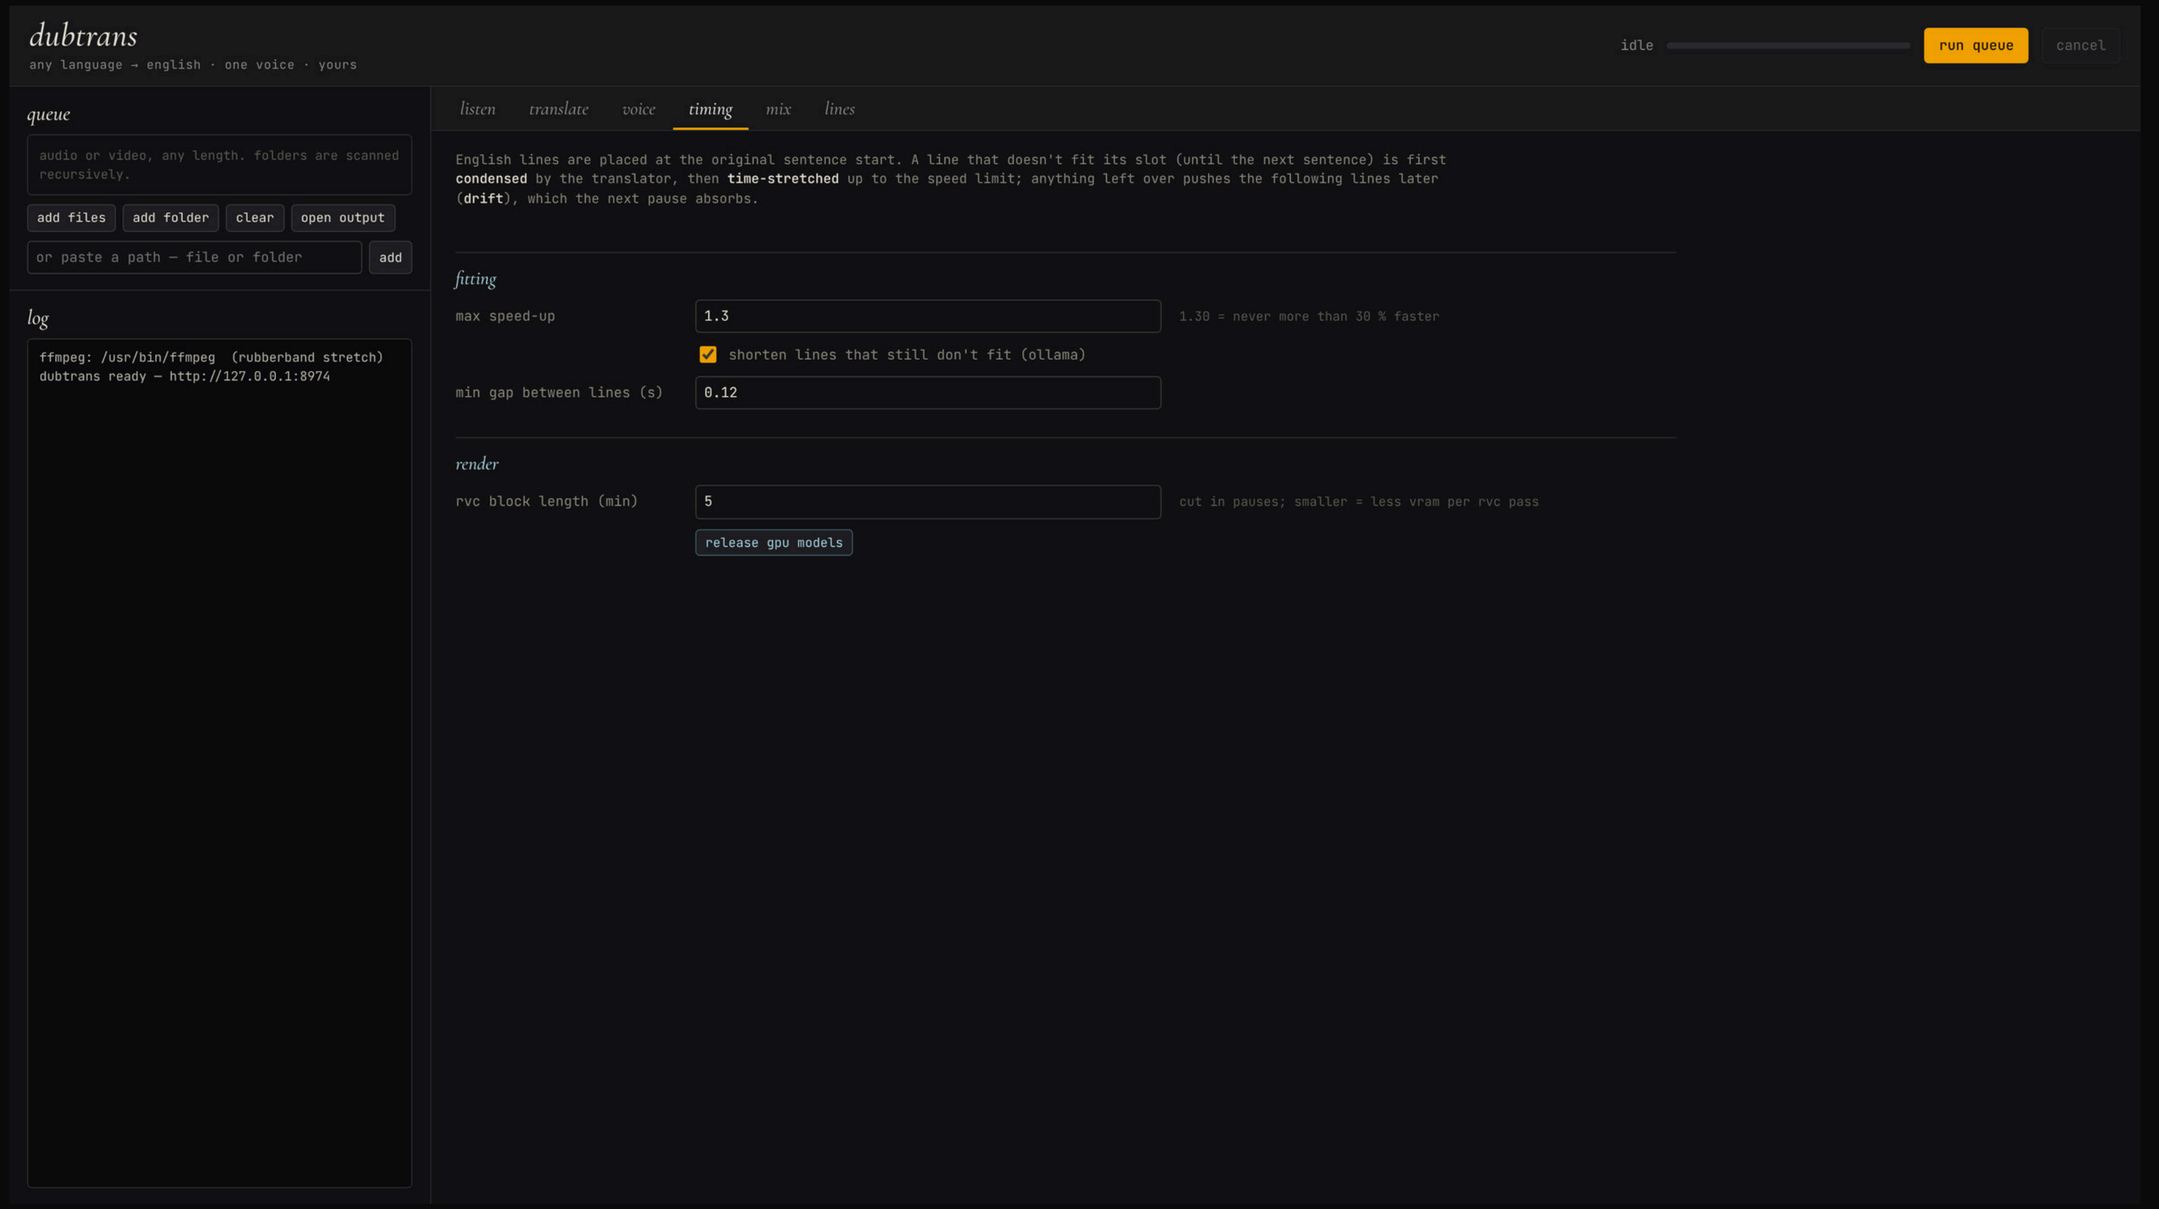
dubtrans

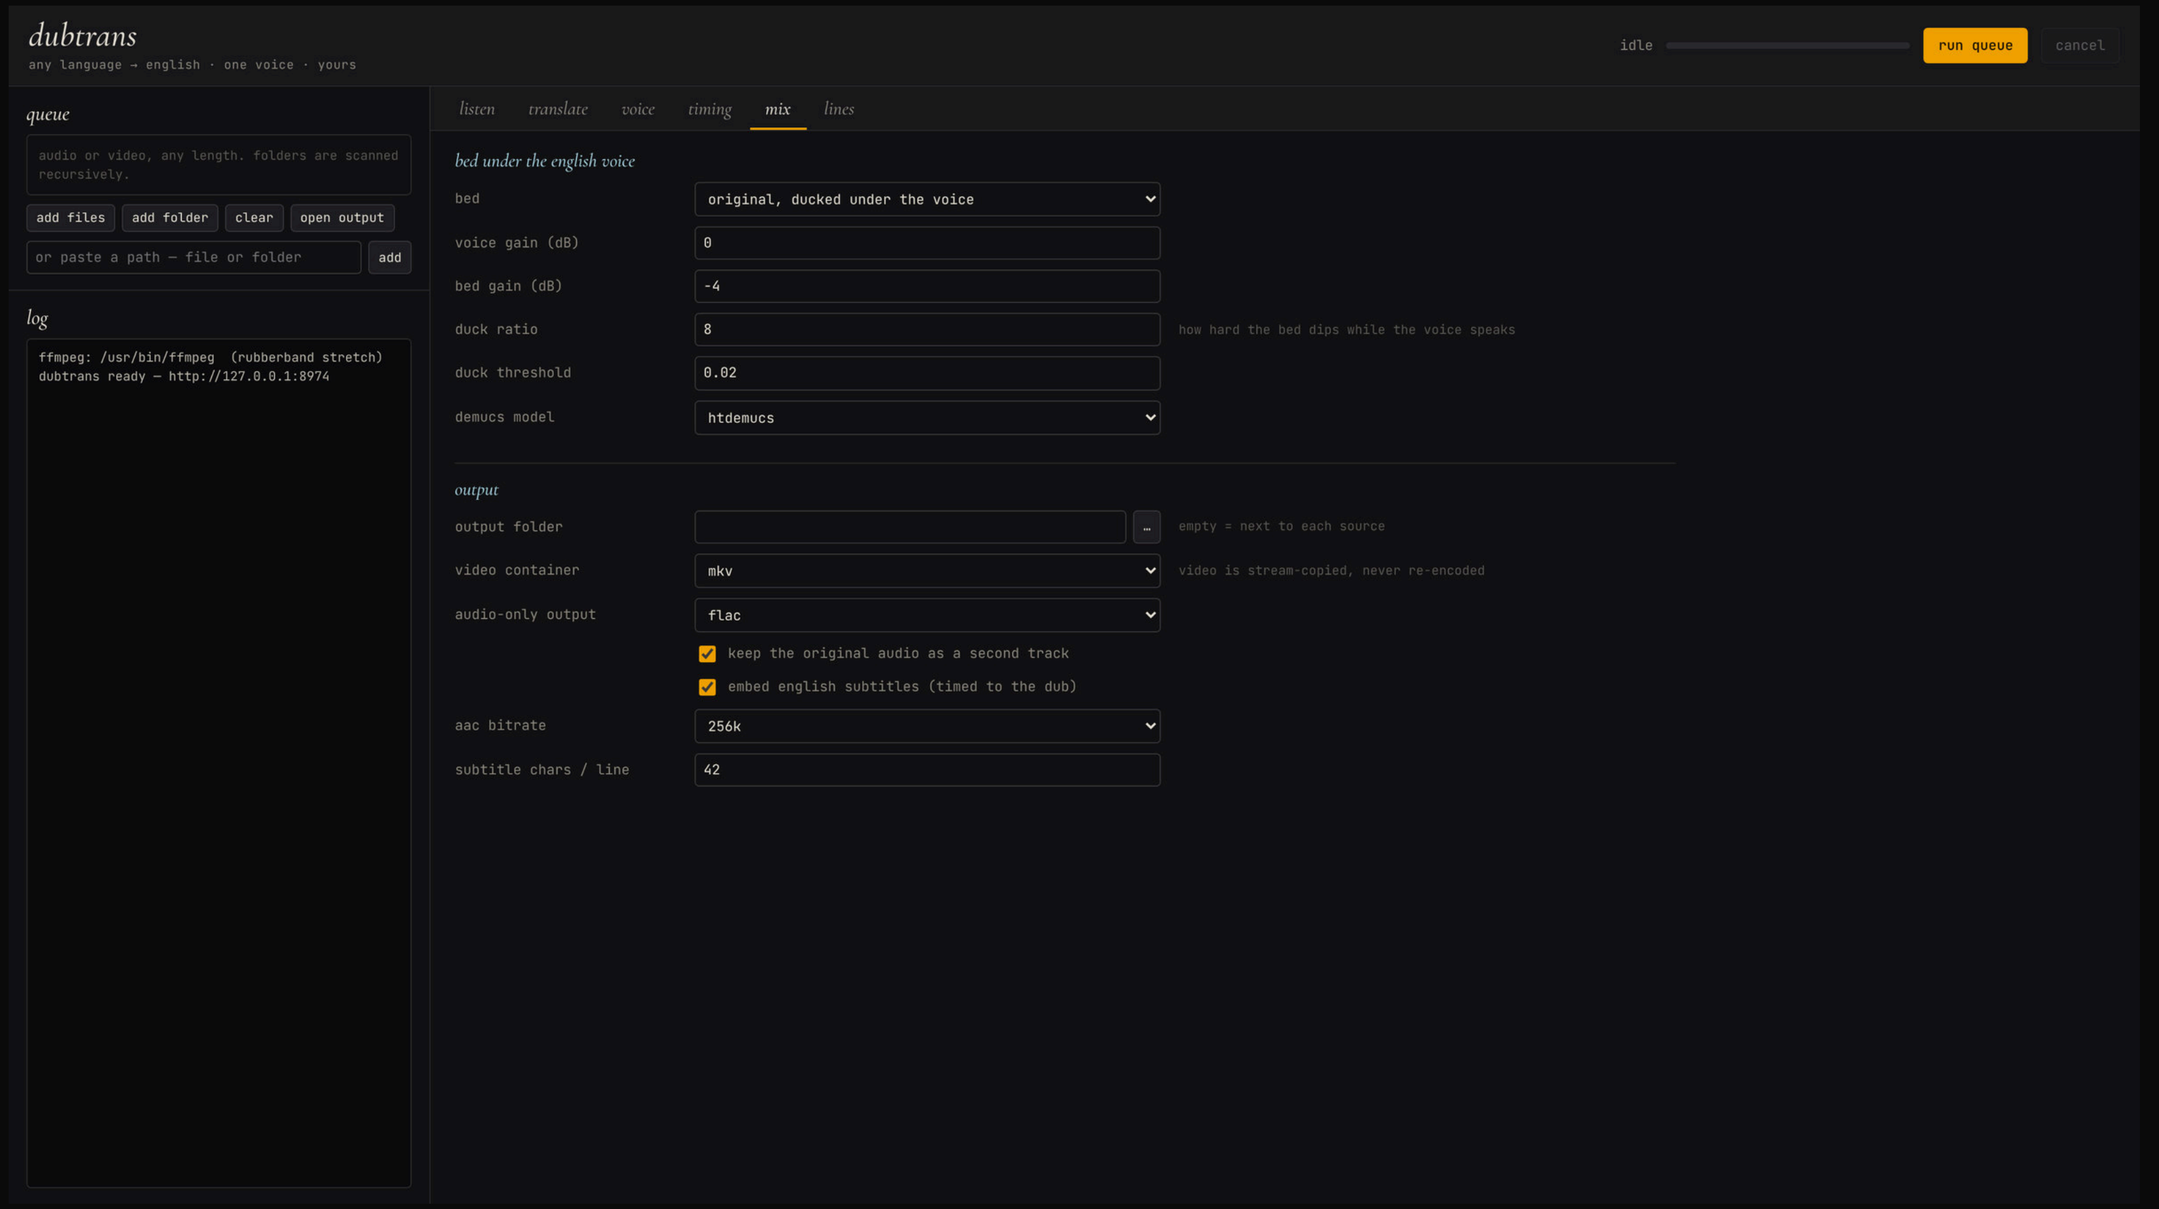
2026

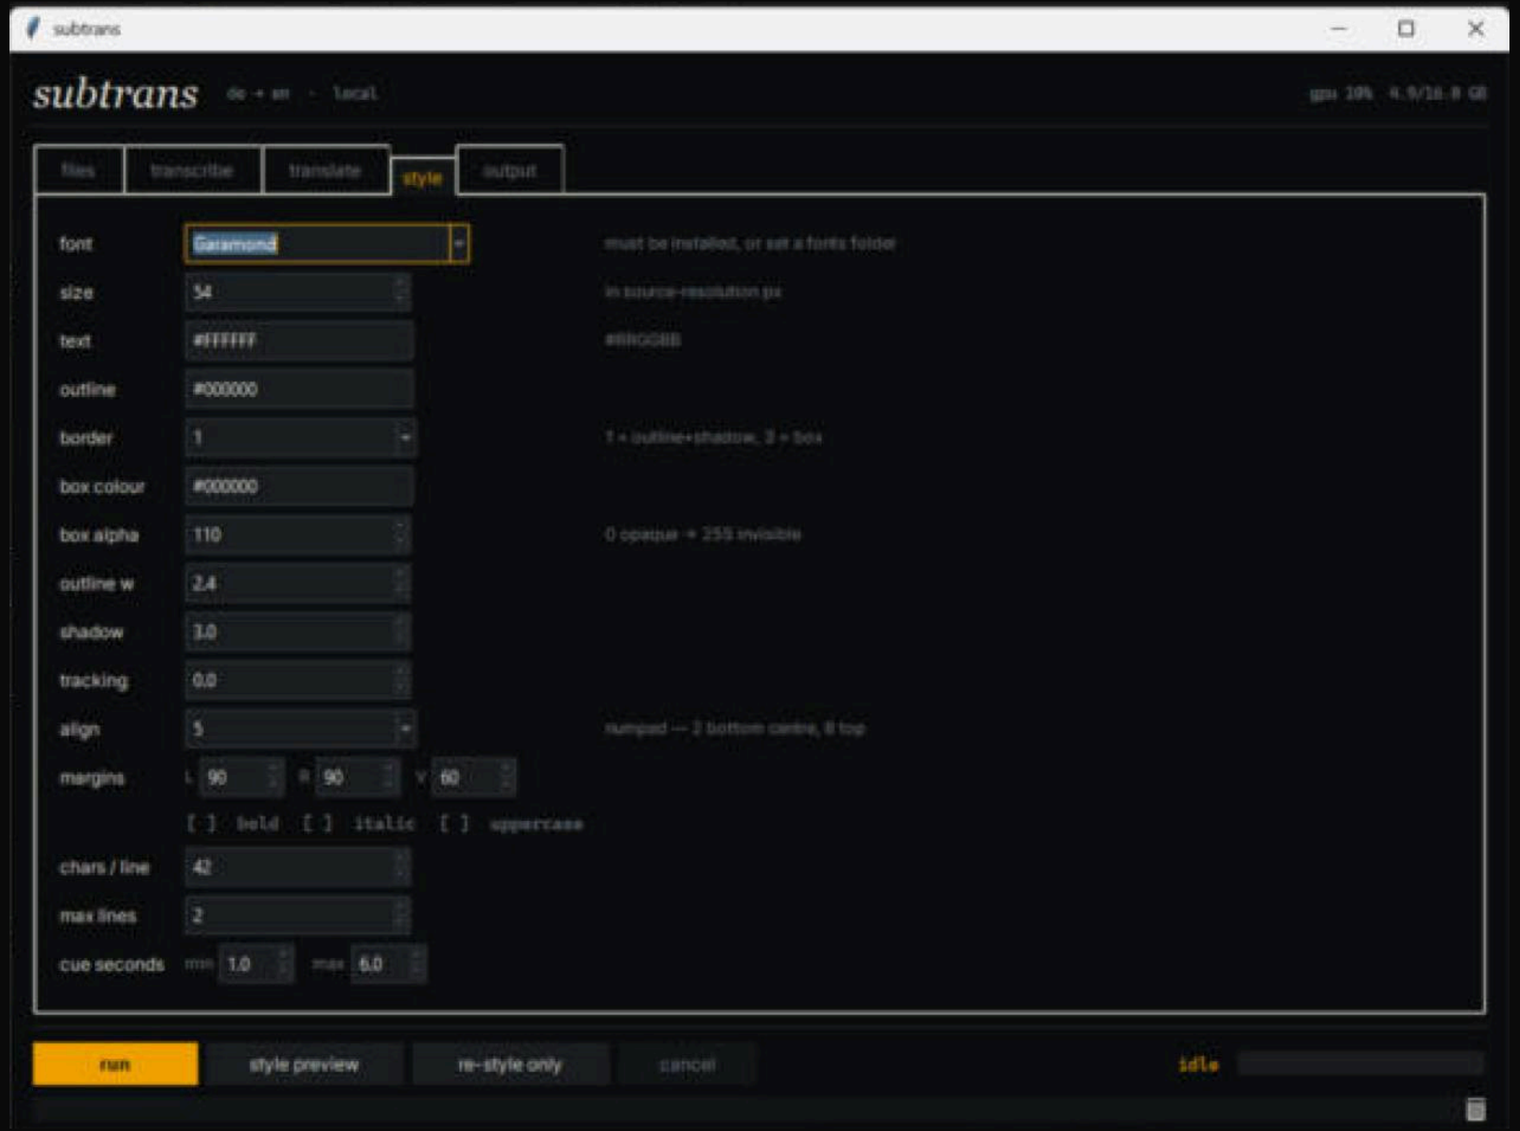
FIELD	Voice and language
TECHNOLOGY	Python, faster-whisper, Ollama (local LLM), Chatterbox TTS, RVC/Applio, Demucs, rubberband, ffmpeg
SOURCE	261018_Dubtrans.zip

Fully local translation dubbing for long recordings. Speech is transcribed (Whisper, language locked after detection), translated by a local language model that is told each line's duration, synthesised, fitted into the original timing by condensation and time-stretching, converted through a trained RVC voice model and mixed over the original or a separated background. Output: video with English dub and original track plus subtitles.


dubtrans
any language → english · one voice · yours
idle
run queue
cancel
queue
audio or video, any length. folders are scanned recursively.
add files
add folder
clear
open output
or paste a path — file or folder
add
log
ffmpeg: /usr/bin/ffmpeg (rubberband stretch)
dubtrans ready — http://127.0.0.1:8974
listen
translate
voice
timing
mix
lines
Every speaker in the source becomes one english voice. Chatterbox speaks the english; if a reference clip is set it clones that voice zero-shot. With RVC on, the whole timeline then goes through your trained Applio model — once per block, so the model loads once, not per line.
speech (chatterbox)
engine
chatterbox
dry = beeps, checks timing + mix without a gpu
reference voice
optional — 10–30 s of the target voice
exaggeration
0.5
0.5 neutral, higher = more dramatic
cfg / pacing
0.5
lower = slower, steadier
temperature
0.8
seed
1234
rvc (your trained model, via applio)
run the dub through the rvc model
applio folder
the folder with core.py — models are read from its logs\
trained model
— nothing in applio\logs —
rescan
model .pth
index .index
optional — falls back to no index if it breaks
pitch (semitones)
0
index rate
0.5
protect consonants
0.33
f0 method
rmvpe
extra applio args
passed through verbatim, e.g. --clean_audio True
applio python
empty = <applio>\env\python.exe
listen
This is a test of the


dubtrans
any language → english · one voice · yours
idle
run queue
cancel
queue
audio or video, any length. folders are scanned recursively.
add files
add folder
clear
open output
or paste a path — file or folder
add
log
ffmpeg: /usr/bin/ffmpeg (rubberband stretch)
dubtrans ready — http://127.0.0.1:8974
listen
translate
voice
timing
mix
lines
English lines are placed at the original sentence start. A line that doesn't fit its slot (until the next sentence) is first condensed by the translator, then time-stretched up to the speed limit; anything left over pushes the following lines later (drift), which the next pause absorbs.
fitting
max speed-up
1.3
1.30 = never more than 30 % faster
shorten lines that still don't fit (ollama)
min gap between lines (s)
0.12
render
rvc block length (min)
5
cut in pauses; smaller = less vram per rvc pass
release gpu models


dubtrans
any language → english · one voice · yours
idle
run queue
cancel
queue
audio or video, any length. folders are scanned recursively.
add files
add folder
clear
open output
or paste a path — file or folder
add
log
ffmpeg: /usr/bin/ffmpeg (rubberband stretch)
dubtrans ready — http://127.0.0.1:8974
listen
translate
voice
timing
mix
lines
bed under the english voice
bed
original, ducked under the voice
voice gain (dB)
0
bed gain (dB)
-4
duck ratio
8
how hard the bed dips while the voice speaks
duck threshold
0.02
demucs model
htdemucs
output
output folder
...
empty = next to each source
video container
mkv
video is stream-copied, never re-encoded
audio-only output
flac
keep the original audio as a second track
embed english subtitles (timed to the dub)
aac bitrate
256k
subtitle chars / line
42


subtrans
subtrans
files
transcribe
translate
style
output
font
Garamond
must be installed, or set a fonts folder
size
54
in source-resolution px
text
#FFFFFF
#RRGGBB
outline
#000000
border
1
box colour
#000000
box alpha
110
0 opaque · 255 invisible
outline w
2.4
shadow
3.0
tracking
0.0
align
5
margins
L 90
R 90
V 60
[ ] bold [ ] italic [ ] uppercase
chars / line
42
max lines
2
cue seconds
min 1.0
max 6.0
run
style preview
re-style only
cancel
idle

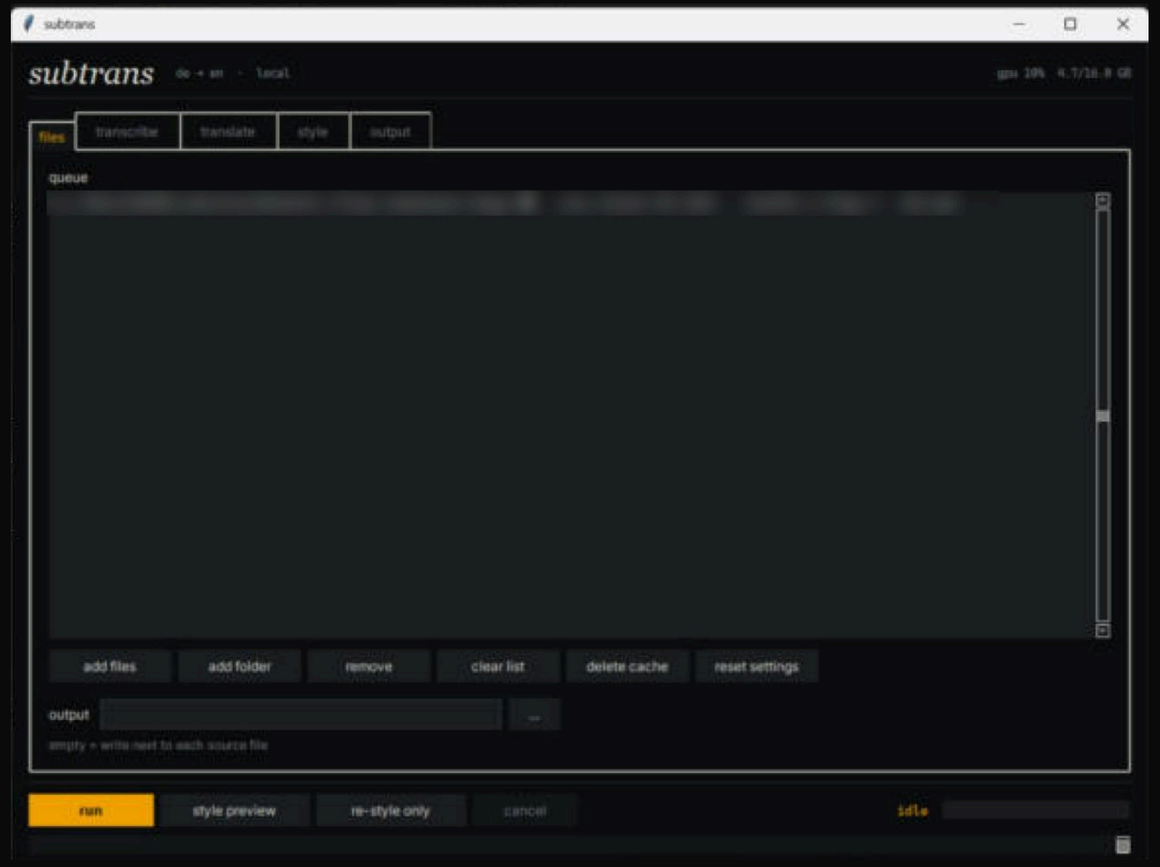
T13

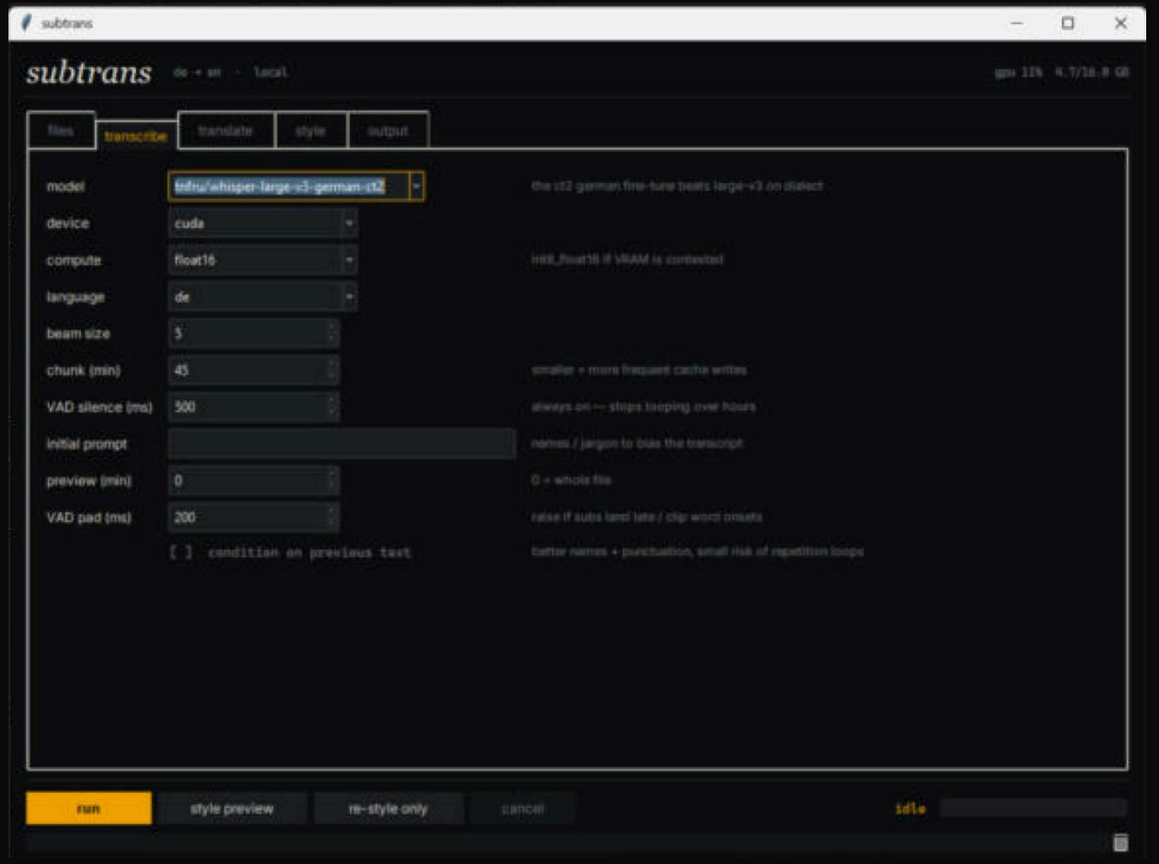
subtrans

2026

FIELD	Voice and language
TECHNOLOGY	Python/tkinter, faster-whisper, Ollama, ASS subtitles, ffmpeg/NVENC
SOURCE	260908_Subtitler.zip

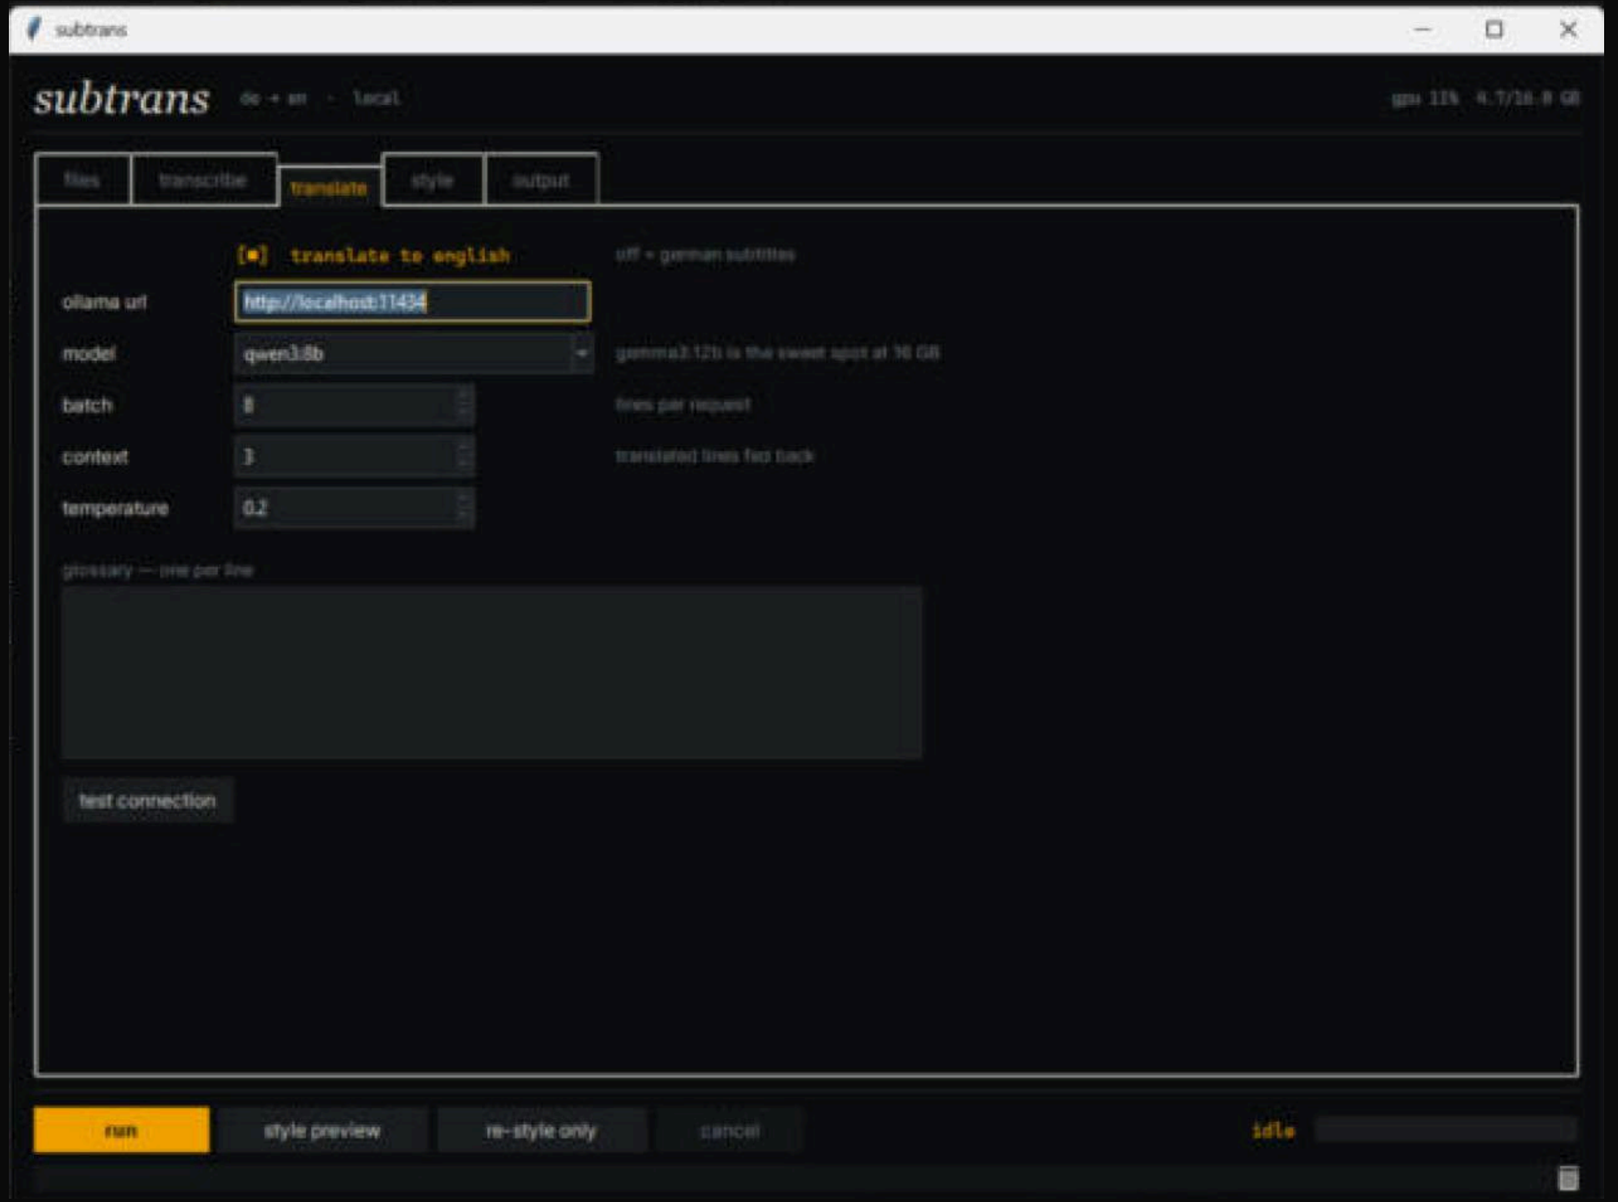
Local subtitling pipeline for recordings of three to seven hours: chunked transcription with voice-activity detection and word timestamps, translation through a local language model, typographic styling as ASS, and either soft muxing or hardware-encoded burn-in. Batch-capable; shares its listening and translation core with dubtrans.

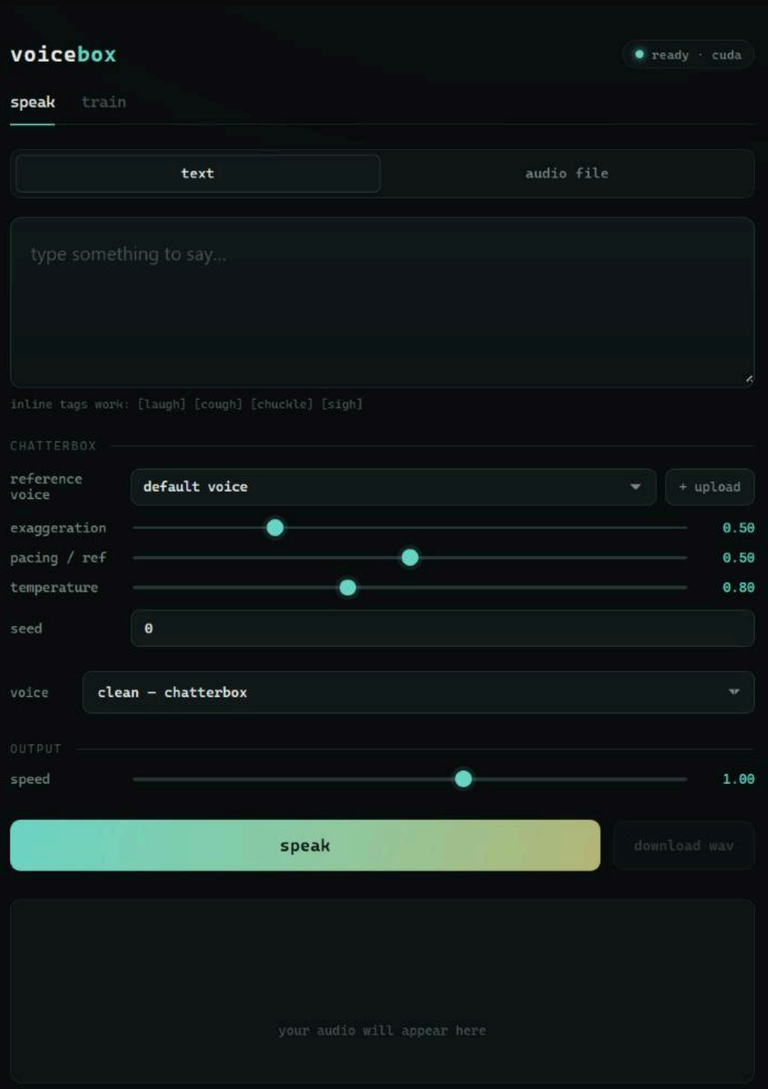




02 FILES · QUEUE 03 TRANSCRIBE

T13 · SUBTRANS


subtrans
subtrans
de → en · local
gpu 11% 4.7/16.0 GB
files
transcribe
translate
style
output
[•] translate to english
off = german subtitles
ollama url
http://localhost:11434
model
qwen3:8b
gemma3:12b is the sweet spot at 16 GB
batch
8
lines per request
context
3
translated lines fed back
temperature
0.2
glossary — one per line
test connection
run
style preview
re-style only
cancel
idle


voicebox
ready · cuda
speak
train
text
audio file
type something to say...
inline tags work: [laugh] [cough] [chuckle] [sigh]
CHATTERBOX
reference voice
default voice
+ upload
exaggeration
0.50
pacing / ref
0.50
temperature
0.80
seed
0
voice
clean — chatterbox
OUTPUT
speed
1.00
speak
download wav
your audio will appear here

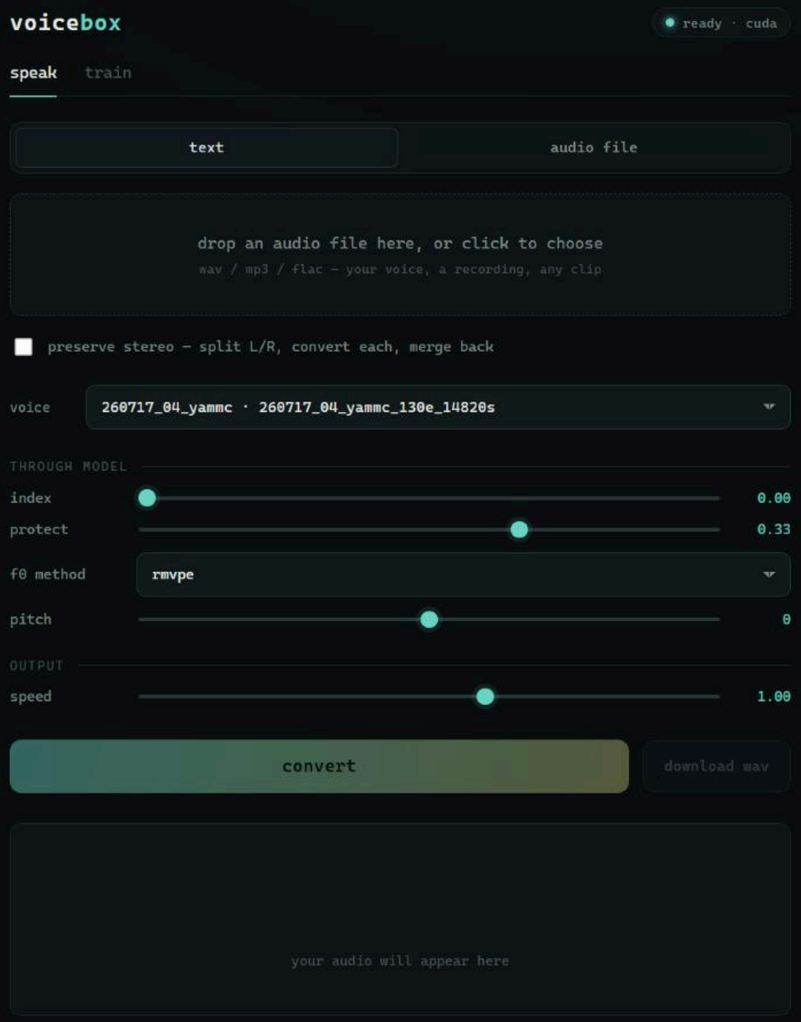
T14

voicebox

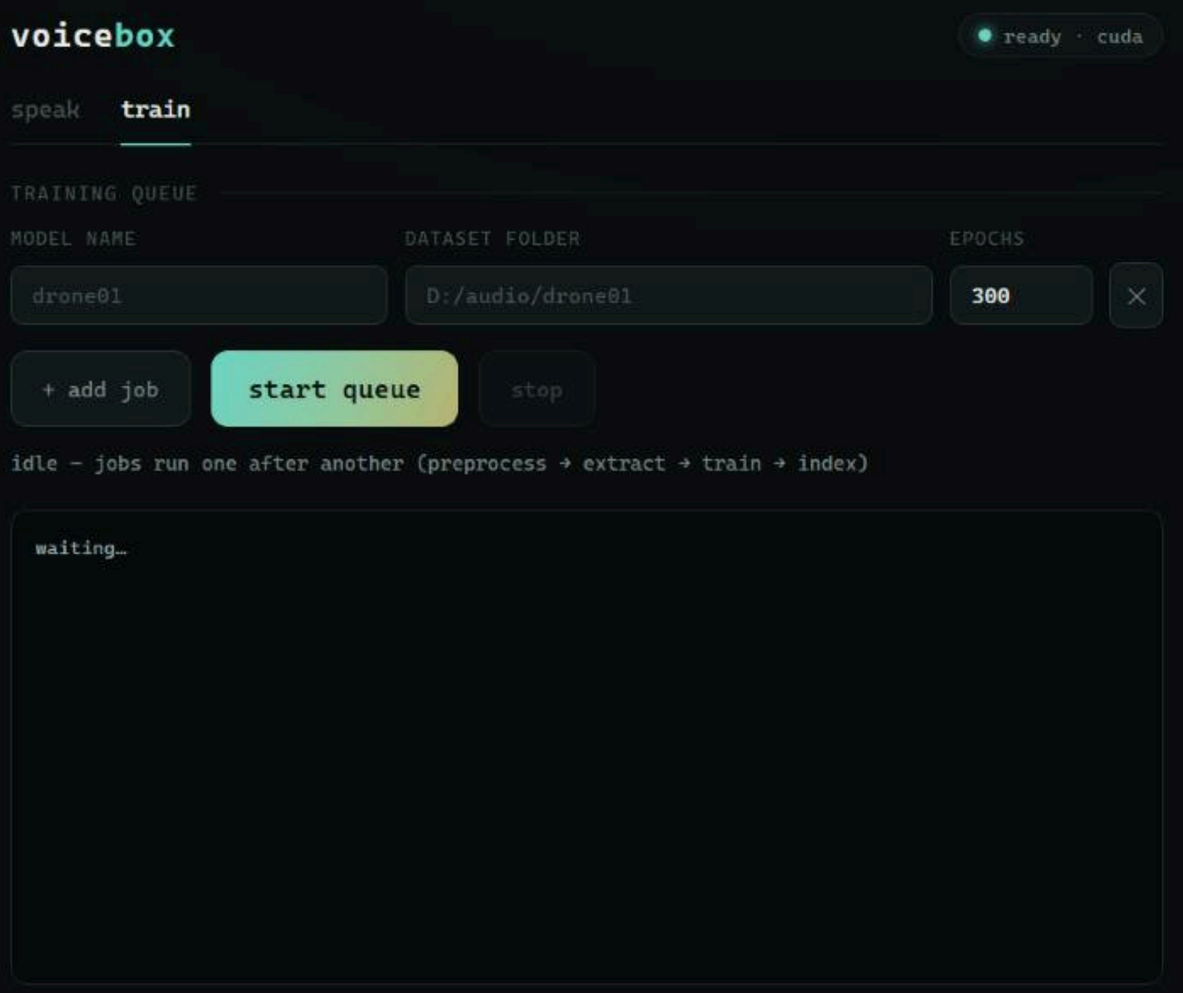
2026

FIELD	Voice and language
TECHNOLOGY	Python, Chatterbox TTS, Applio/RVC, aiohttp web UI
SOURCE	260716_VoiceCloner.zip

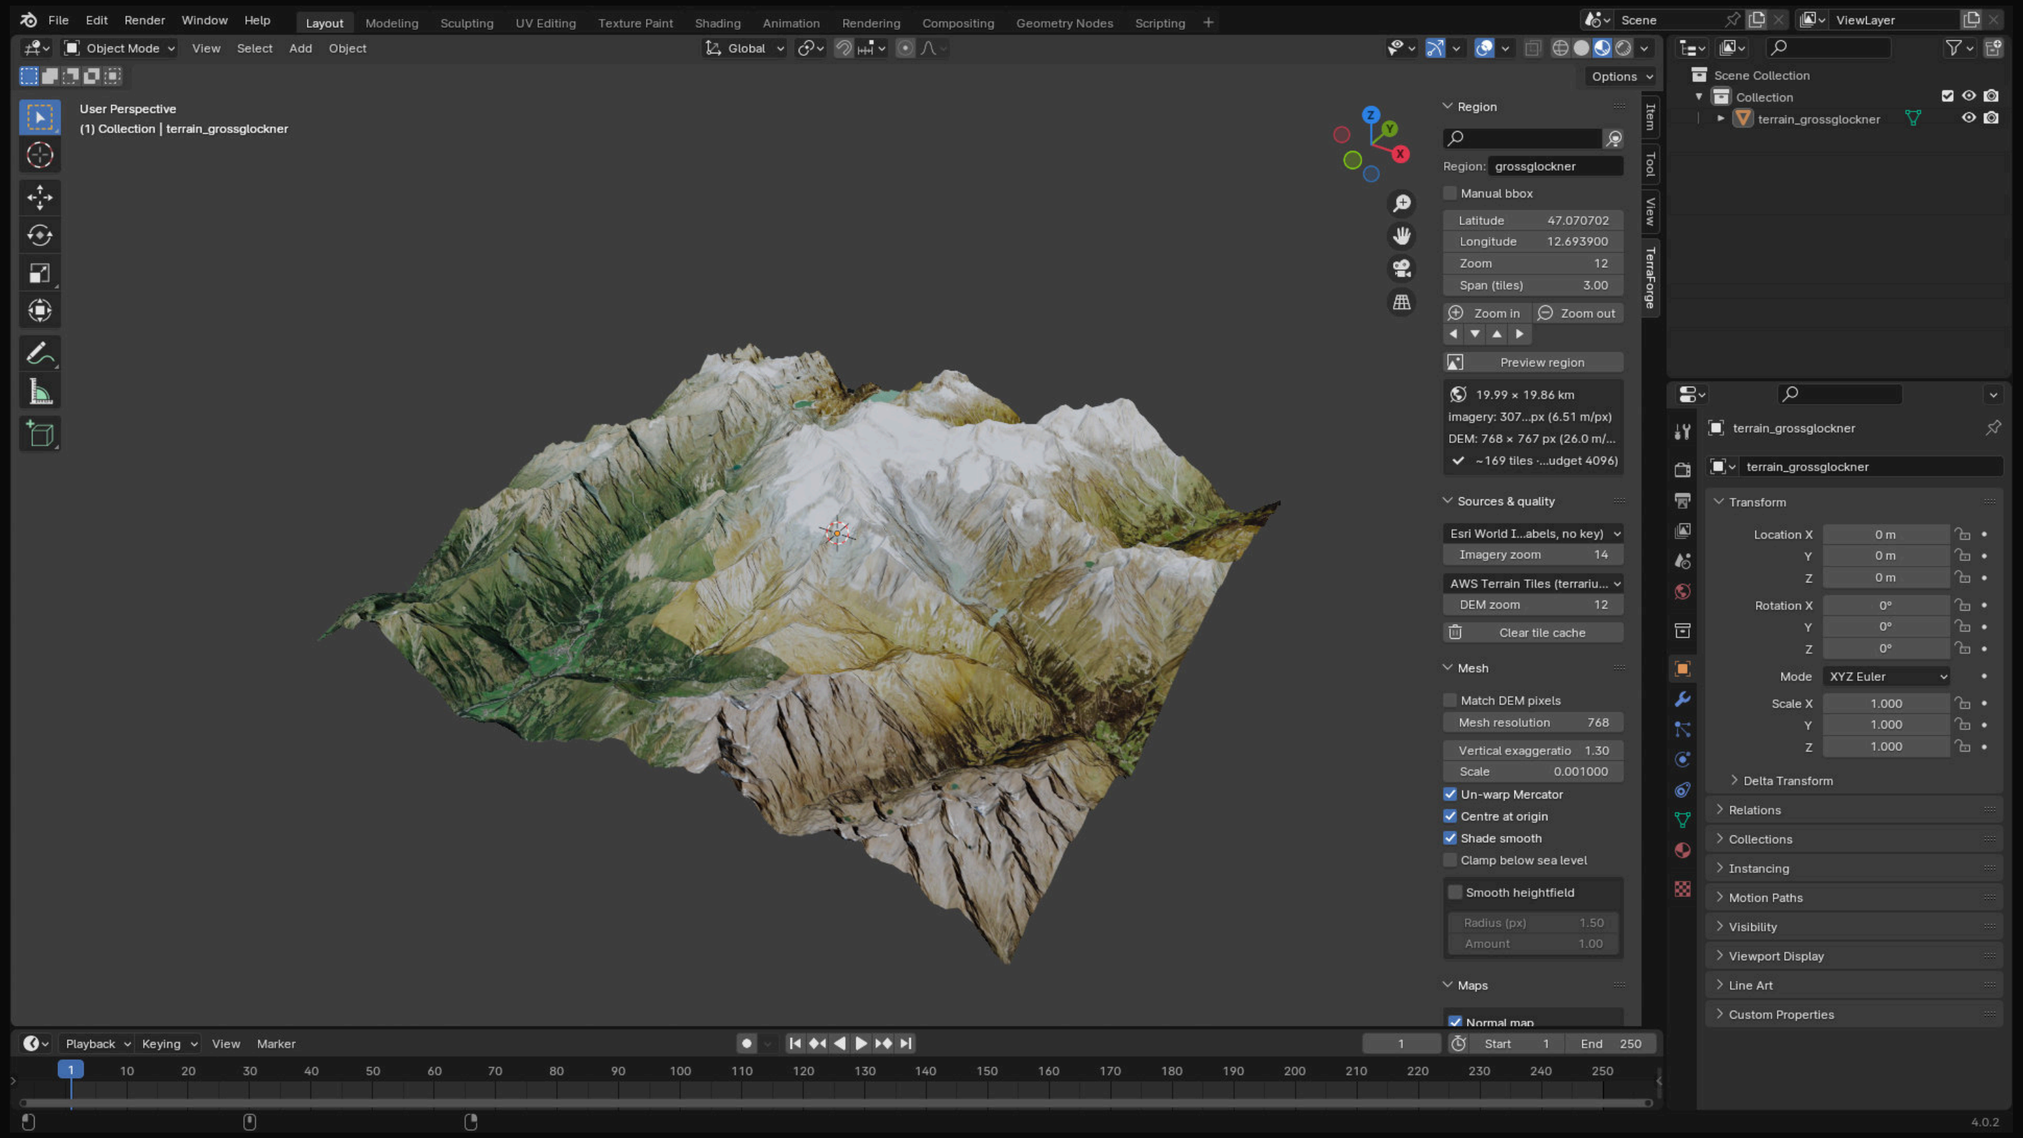
A local text-to-speech environment whose voice can be a model trained on arbitrary, including non-vocal, audio. Typed text or any input file is synthesised and passed through a trained RVC model; reference voices, expressive tags, extreme slow-down and stereo-preserving conversion are available, and training jobs can be queued from the same interface.



02 CONVERT · TRAINED VOICE



03 TRAIN · QUEUE



TerraForge

2026

SPACE, TERRAIN AND BODY

01 · Blender · TerraForge panel

T15

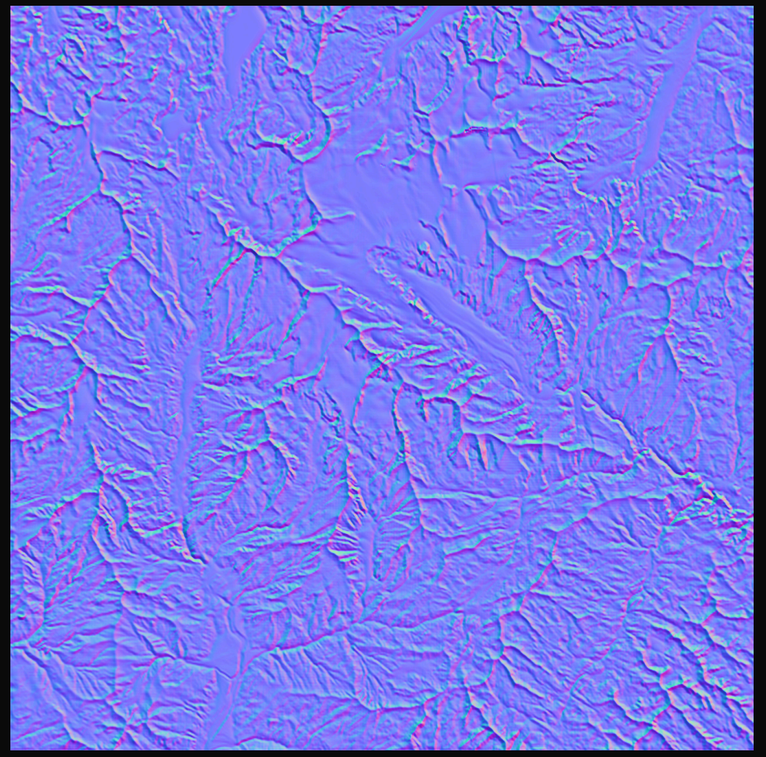
TerraForge

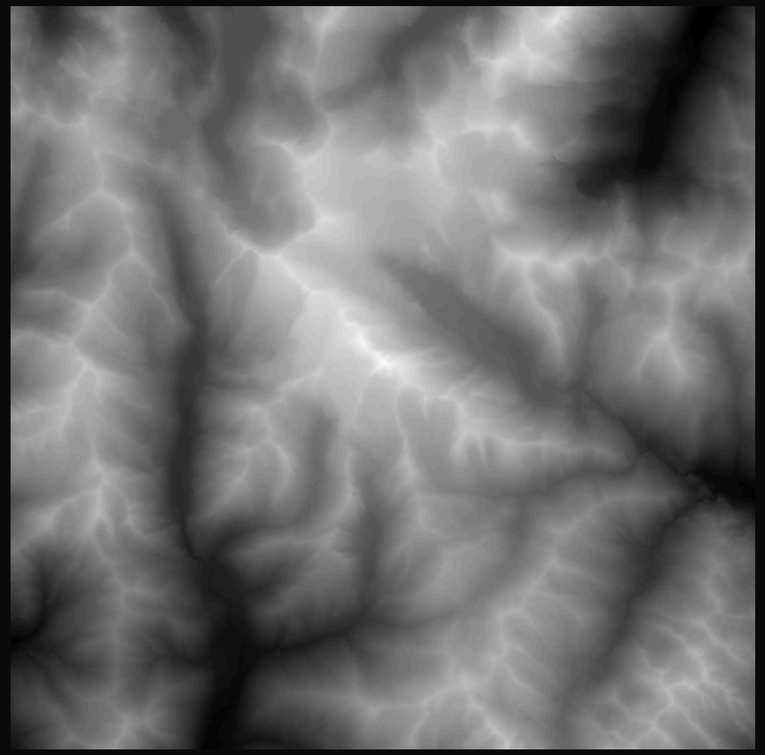
2026

FIELD	Space, terrain and body — 3D and game-engine tools
VERSION	1.0
TECHNOLOGY	Blender add-on (Python/numpy), open elevation tiles (Terrarium), Esri/USGS/Mapbox imagery
SOURCE	terraforge.py

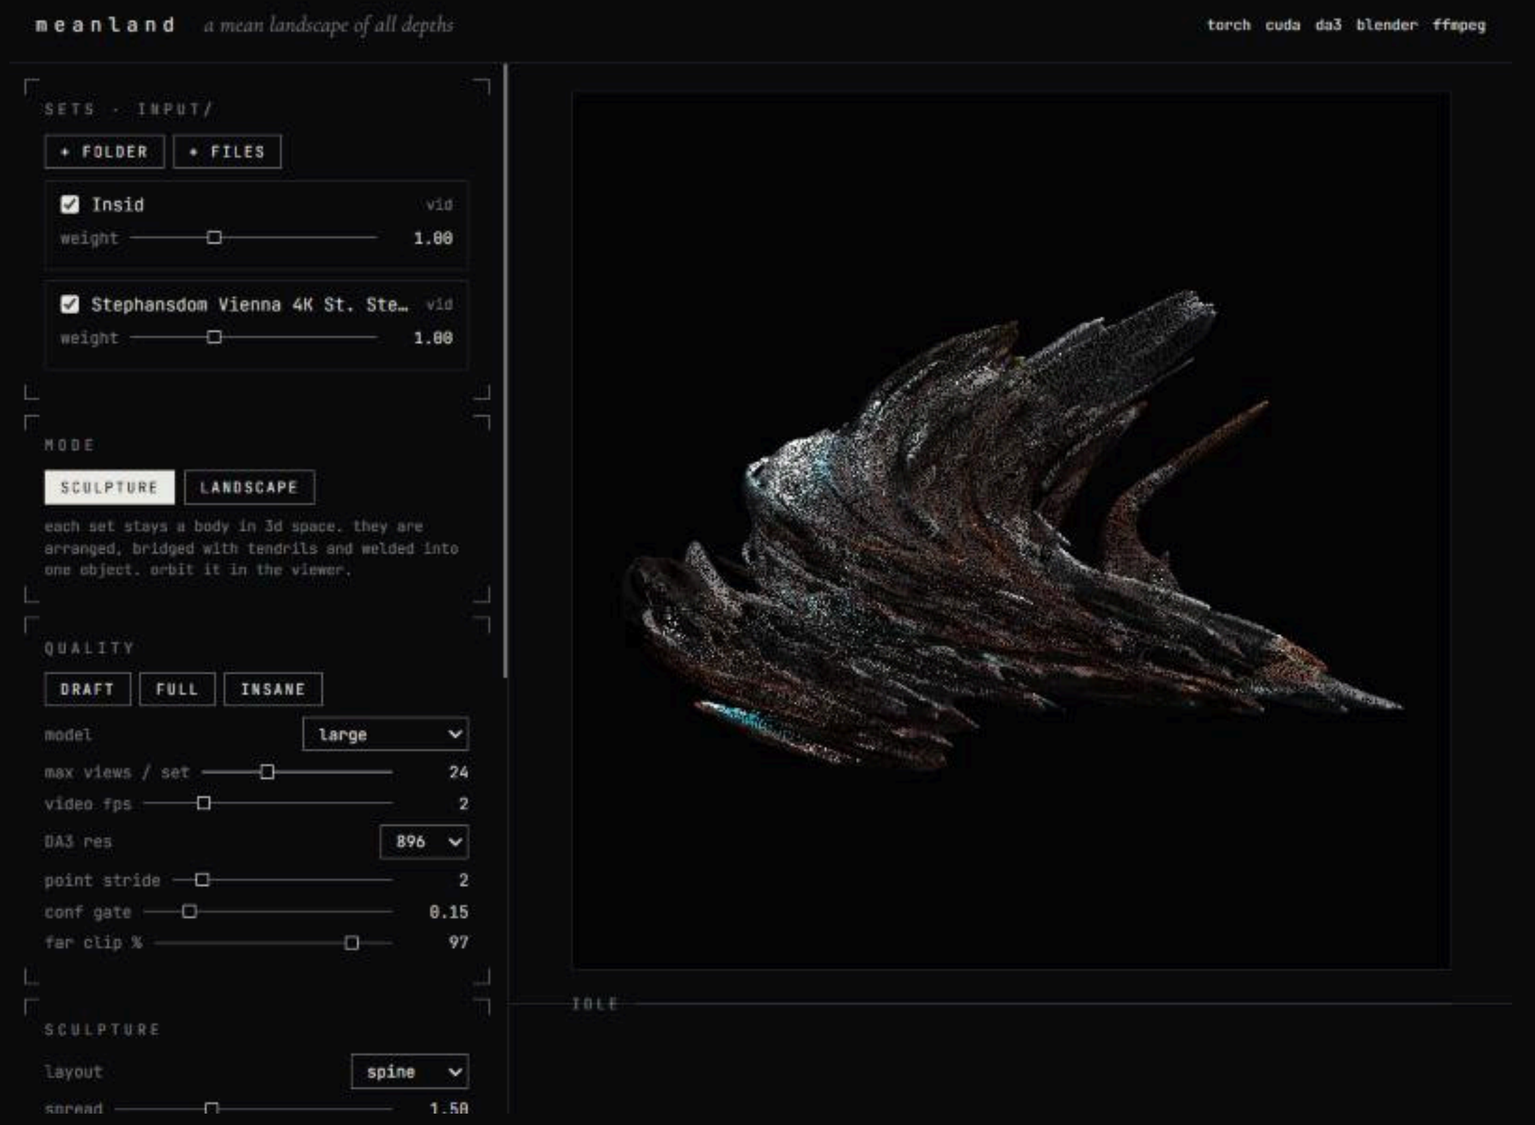
A Blender add-on that fetches elevation and satellite imagery for any region on earth, builds a metrically correct (un-warped Mercator) terrain mesh and derives normal, 16-bit height and ambient-occlusion maps, exported with dated textures as FBX/OBJ for game engines. Plate: Großglockner, 20 × 20 km.







02 ALBEDO · ESRI Z14 03 NORMAL MAP 04 16-BIT HEIGHT


meanland
a mean landscape of all depths
torch cuda da3 blender ffmpeg
SETS · INPUT/
+ FOLDER
+ FILES
Insid
vid
weight
1.00
Stephansdom Vienna 4K St. Ste…
vid
weight
1.00
MODE
SCULPTURE
LANDSCAPE
each set stays a body in 3d space. they are arranged, bridged with tendrils and welded into one object. orbit it in the viewer.
QUALITY
DRAFT
FULL
INSANE
model
large
max views / set
24
video fps
2
DA3 res
896
point stride
2
conf gate
0.15
far clip %
97
SCULPTURE
layout
spine
IDLE

T16

meanland

2026

FIELD	Space, terrain and body — 3D and game-engine tools
VERSION	v5
TECHNOLOGY	Python, Depth Anything 3 (CUDA), browser UI, FBX export
SOURCE	meanland_v5.zip

Fuses several image sets or videos through Depth Anything 3 (depth and camera pose) into an averaged landscape heightfield or a free-standing sculpture, exported as textured FBX for Unity, and as per-frame point clouds for 4D sequences.



meanland *a mean landscape of all depths*

torch cuda da3 blender ffmpeg

WELD

backend: weld

weld k · flesh: 0.10

shell · crust: 0.035

silk · anti-spike: 1.0

smoothing passes: 12

detail · sculpt res: 320

poisson depth: 9

trim: 0.04

CONNECT

one continuous form: ☑

tendrils / pair: 24

tendril thickness: 0.020

tendril wander: 0.15

keep fragments: ☐

DEFORM

twist °: 25

bend: 0.12

stretch: 1.00

BATCH

variations: 1

seed: 0

OBJECT

world size (unity u): 100

texture res: 2048

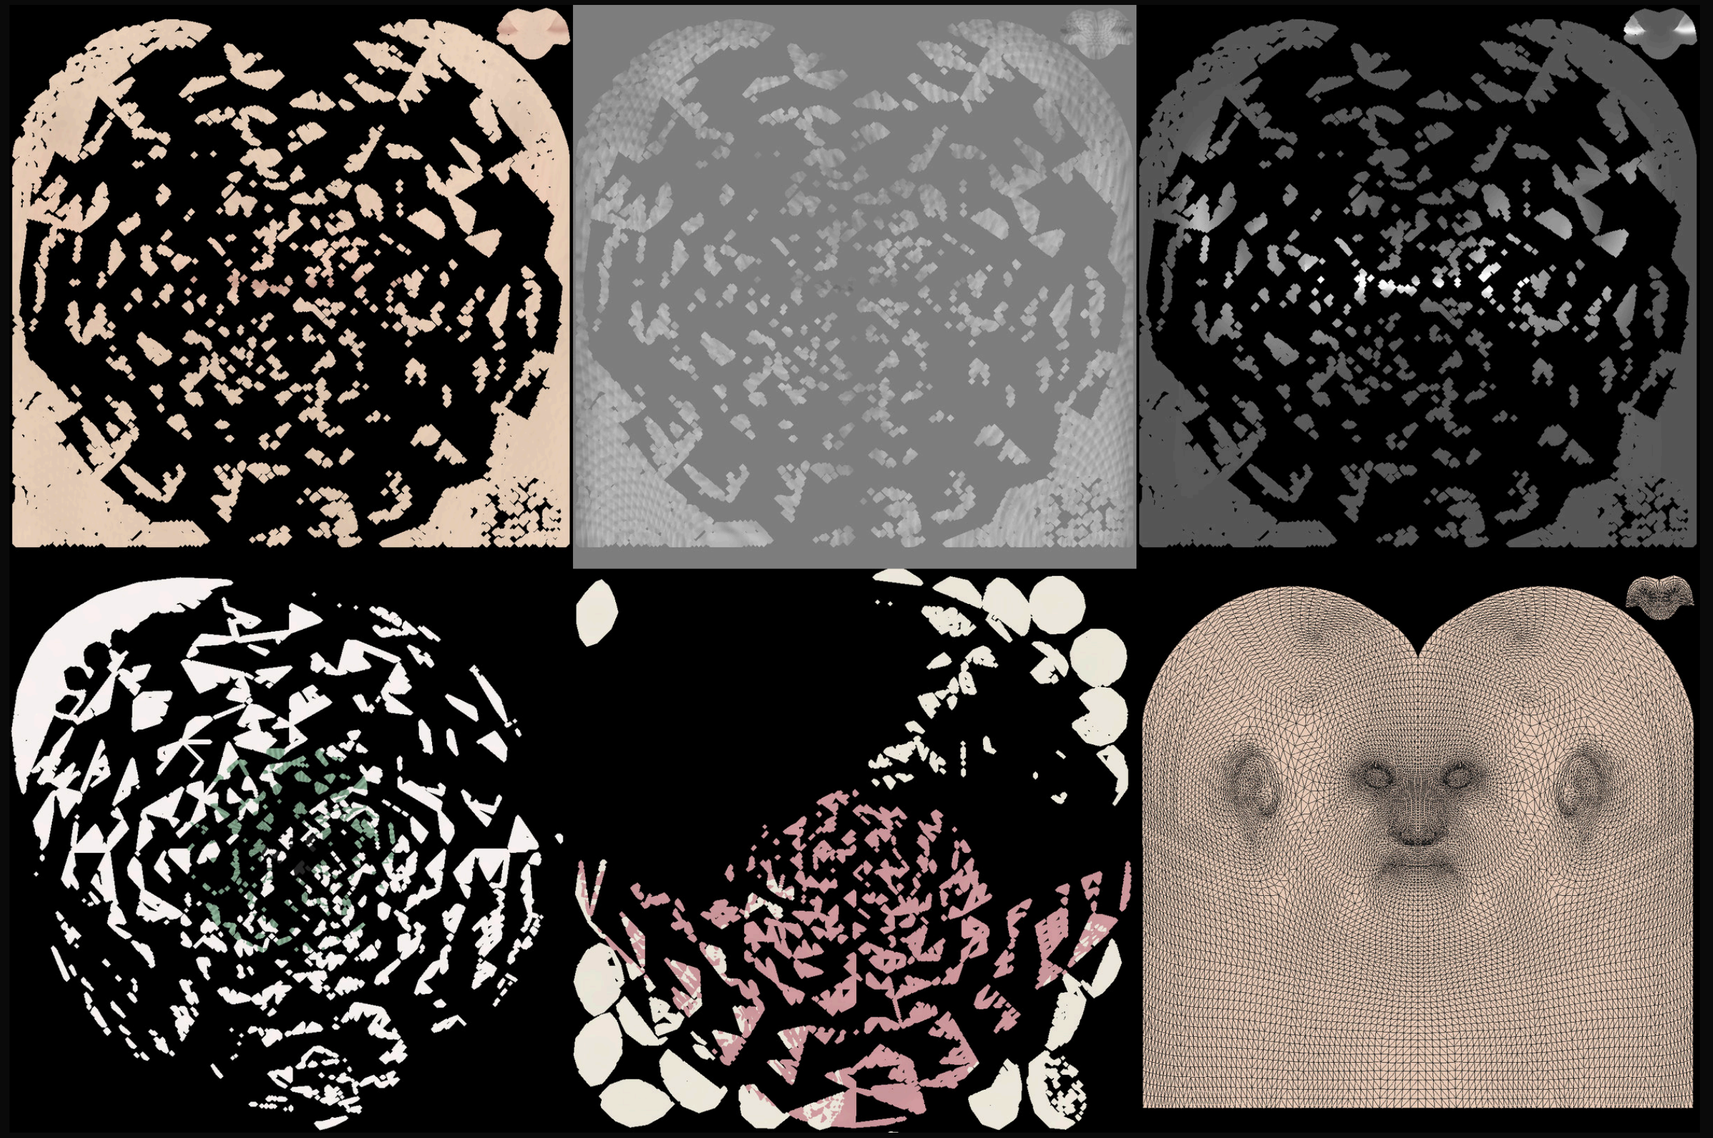
export FBX (blender): ☑



IDLE

02 FUSED SCULPTURE · PREVIEW

03 WELD · CONNECT · DEFORM



GNM head pipeline

2026

SPACE, TERRAIN AND BODY

01 · texture kit · bake sheets

T17

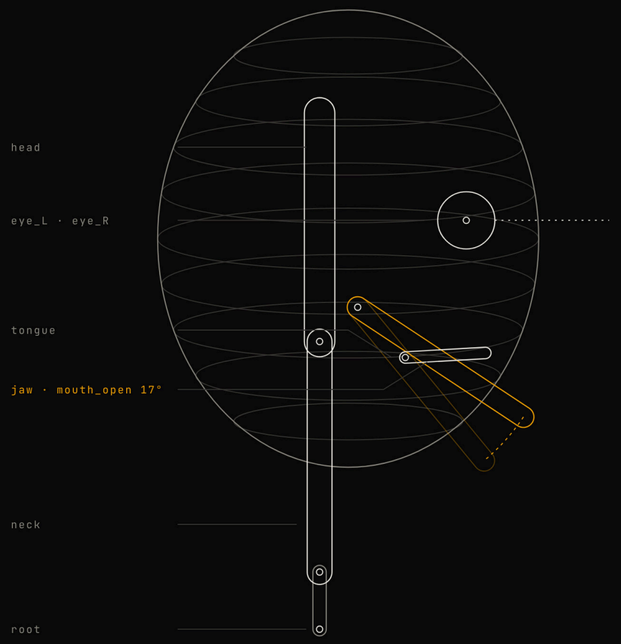
T17

GNM head pipeline

2026

FIELD	Space, terrain and body — 3D and game-engine tools
VERSION	rig v0.6 · texture kit v1.0 · projector 2
TECHNOLOGY	Blender add-ons (Python/numpy), mediapipe/manual landmarks, thin-plate-spline warp; Unity HDRP
SOURCE	gnm_unity_rig.py · 260722_GNM_textureKit.zip · 260722_GNM_textureKit_projector2.zip

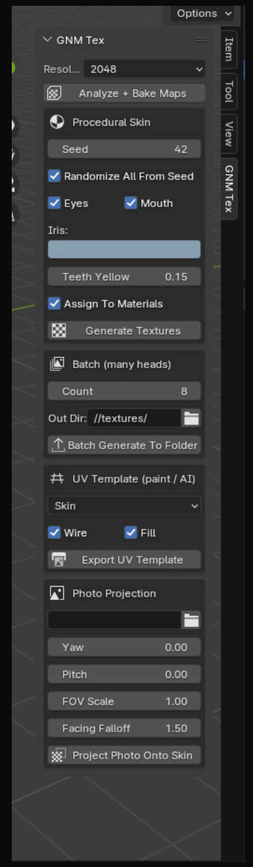
A pipeline for neural-generated (GNM) heads. The rig add-on classifies regions from material slots, solves the anatomical frame from eye and teeth centroids and builds a reproducible Root/Neck/Head/Jaw/Tongue/Eye rig with baked viseme and blink shape keys for Unity. The texture kit bakes position and normal maps per UV region and synthesises seed-driven skin, eyes and mouth; the projector conforms any frontal (e.g. AI-generated) face photo onto the head via landmark alignment and thin-plate-spline warping.



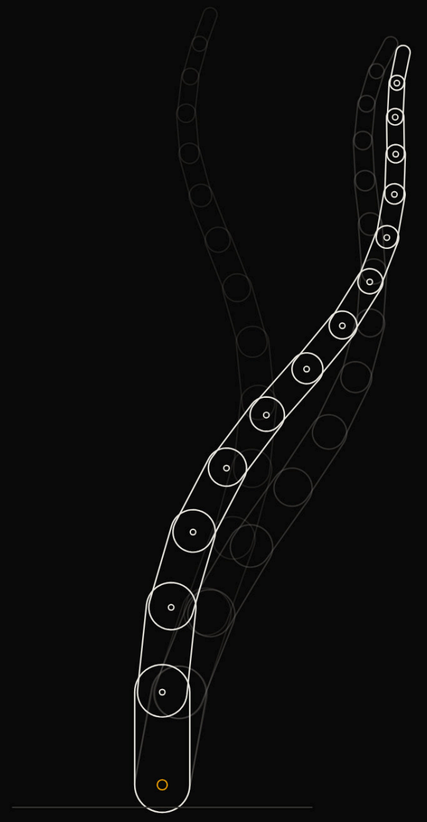
rig solved from material slots · same frame on every generated head



9 visemes baked as shape keys from three axes: open · wide · round



02 RIG (ILLUSTRATION) 03 VISEME SHAPE KEYS (ILLUSTRATION) 04 BLENDER · GNM TEX PANEL



tentacle · 14 bodies · 13 joints

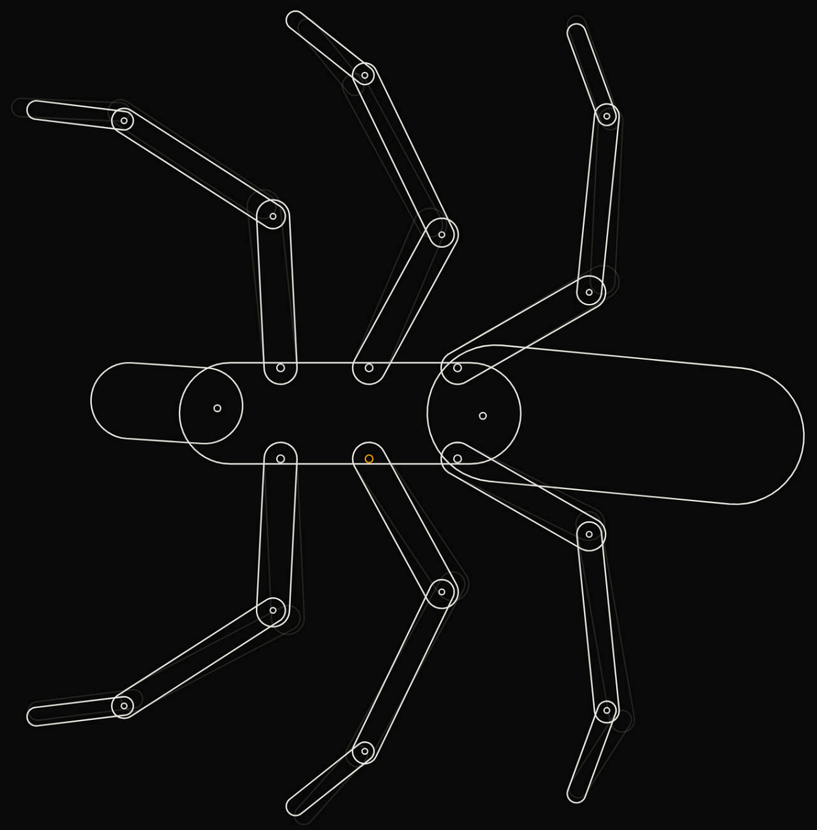
T18

AutoRagdoll

2026

FIELD	Space, terrain and body — 3D and game-engine tools
VERSION	Unity 6.5
TECHNOLOGY	Unity 6.5 HDRP, C# (runtime + editor), PhysX
SOURCE	AutoRagdoll_Unity65.zip

A rig-agnostic ragdoll generator for Unity: colliders and joints are fitted from skin weights through the bind pose, so bipeds, quadrupeds, insects, tentacles and mechanical rigs work without humanoid mapping; unrigged meshes are chunked into bodies. A driver adds eight motion layers, 26 presets and audio input.



hexapod · 21 bodies · no humanoid map



unrigged mesh · 18 chunks · fitted joints

02 HEXAPOD (ILLUSTRATION)

03 UNRIGGED MESH (ILLUSTRATION)



Resonode

2026

PERFORMANCE CONTROL

01 · preset · four layers

T19

T19

Resonode

2026

FIELD	Performance control
TECHNOLOGY	Node.js OSC bridge, HTML interface; Resolume Avenue 6
SOURCE	Resonode.zip

An OSC control surface for Resolume built for slow performance: long fades, few gestures, nothing cutting that should glide. The fade engine runs in the bridge rather than the browser, so a twelve-minute crossfade keeps its exact curve even if the interface is closed; parameters, soft triggers, clip grids and slow sequences are arranged as nodes, with cues and remote control.


RESONODE
bridge
to 127.0.0.1
7000
from 7001
send 30 Hz
time 1.00×
colours spectrum
Stop all · 5
Black
Float
Pin
Presets
four layers
four layers
Load
name
Save
keep current IP & ports when loading
Export
Import
Cues
Capture
01 black wait
02 layer 1 alone wait
03 one and two 30s
04 everything up wait
GO space
Stop
black — waiting for GO
Remote
Let another machine drive this one: point Max, Reaktor or another sender at the "from" port and map an address here.
master
MST
3.2s
float
/composition/master
0.000
TRIGGER
layer 1 → next clip
L1
C1
soft
/composition/layers/1/clips/1/connect
dip a fader
layer 1
FADE INTO IT
out 12 sineInOut hold 0.3
in 12 sineInOut from 0 to 1
...and out again
clips
layers 1 – 4 clips 1 – 8
fade the clip in over 6 s
L1
L2
L3
L4
layers
end · fade 20.0s go
layer 1
L1
15.2s → layer 2
/composition/layers/1/video/opacity
hand · cut +2s go
layer 2
L2
15.2s → layer 3
/composition/layers/2/video/opacity
layer 3
L3
15.2s
/composition/layers/3/video/opacity
layer 4
L4
15.2s
/composition/layers/4/video/opacity
TRIGGER
to 1 in 20 s sineInOut
slow sequence
START
6 clips
stop when faded
hold 180 s · crossfade 90 s equal power
+ clip
one pass ≈ 21:00 + 2 waiting for you
Send
5
/composition/…
1
float
Send
Check every address
Asks Resolume about every address on the canvas and marks the ones it doesn't answer on. Needs OSC Output to port 7001. For an exact address: Shortcuts → Edit OSC in Resolume, then click the control.
Incoming
cue → black
listening for OSC on 7001
loaded preset "four layers"
OSC out → 127.0.0.1:7000 · rate 30 Hz
+ parameter
+ soft trigger
+ clip grid
+ clip chain
+ frame
Tidy
Fit
100%
drag a header to move · drag a right-hand port onto another node to hand off · a add · f fit

START
What it is
Install · Windows
Install · macOS
Setting up Resolume
Keeping it in front

THE SURFACE
The canvas
Colours
Parameters
Seeing time pass

PLAYING CLIPS
Clip grid
Soft triggers
Clip chain
An installation

STRUCTURE
Cues
Cables
Remote control
Presets & files

REFERENCE
OSC addresses
Keyboard
If something is wrong

Resonode

An OSC control surface for Resolume Avenue 6, built for slow work: long fades, few gestures, nothing cutting when it should glide.

What it is

Two files do the work. `bridge.js` is a small server that talks OSC to Resolume and serves the interface. `ui.html` is the interface, which you use in a browser. There is nothing to install beyond Node.js — no npm packages, no internet connection needed at the venue.

The one thing worth understanding: the fade engine runs in the bridge, not in the browser.

Browsers throttle timers in background tabs, which would visibly wreck a ten-minute fade. Here the browser is only a view. A twelve-minute crossfade keeps its exact shape whether the window is in front, hidden, or closed entirely. You can shut the browser mid-fade and the fade finishes. Open it again and you rejoin a show already in progress. Two devices — laptop and phone — stay in sync on the same state.

Install and run · Windows

`Windows 10 / 11`

1. Install **Node.js** from nodejs.org — the LTS build. Accept the defaults.
2. Put the `resonode` folder somewhere sensible (Documents is fine — not inside a zip, and not in Program Files).
3. Double-click **start.bat**.

A black console window appears and stays open — that is the engine; leave it running. The interface opens automatically at `http://127.0.0.1:8962`, in a chromeless Chrome or Edge window if you have one.

To stop: close the interface window, then press `Ctrl` + `C` in the console, or just close the console.

If Windows Defender asks about the firewall

Allow it on private networks. Sending to Resolume works either way, but blocking it stops OSC *coming back* from Resolume, which quietly disables clip names, playing-clip highlighting and the address checker.

If the port is busy

Open a command prompt in the folder and start it on another port:

```
node bridge.js --ui-port 8963
```

Thomas Wagensommerer

VIENNA