


HumanizerHost

Make your virtual instruments breathe — automatic MIDI CC expression
for long notes and whole phrases



v0.8.0 · VST3 instrument · macOS & Windows · built for Cubase, Dorico and other DAWs

What is HumanizerHost?

A VST3 instrument that hosts your libraries

Load up to sixteen VST3 instruments (Kontakt, Synchron Player, BBC SO, SWAM, ...), layered and summed, each with an optional insert FX and its own volume — rendered in parallel on your performance cores, inside the audio thread's workgroup (one shared worker pool, however many plugin instances the session uses).

It writes the expression for you

Incoming MIDI is analyzed and the plugin automates your instrument's dynamics CCs (CC1, CC11, ...) — so sustained notes swell, settle and fade like a real player.


Two layers of humanization

Phrase humanization shapes the dynamics across whole musical phrases; long-note curves shape every sustained note individually.

Everything travels with your session

Hosted plugins, their full state, curves and settings are saved in the DAW project. Export/render is identical to what you hear.


The instrument rack



Each of the 16 slots receives its own rendition of the humanized MIDI. Hosted editors open in floating windows.

The trick: it listens ahead


All MIDI reaches the hosted instruments delayed by the analysis window — so when a note plays, the plugin already knows its length.



One knob, two jobs — window = delay (0.5–10 s, default 5 s). Playback-only mode (default: on) keeps live input instant while the transport is stopped — the humanizer engages during playback and export only.

The delay is deliberate and identical on the timeline and when exporting: what you hear is exactly what renders.

Long-note humanization — before & after



Notes are never modified — only the CC lanes are written. One curve owns a CC at a time: the most recent long note wins.

The Notes tab — where you shape it

The screenshot shows the HumanizerHost software interface, specifically the Notes tab. The top bar includes the HumanizerHost logo, version v0.8.0, copyright © 2026 Franck Dansaert, and an 'Enabled' status button. Below the top bar are four tabs: Phrase, Notes (selected), Instrument, and Monitor. The main interface is divided into several sections. At the top, there are three sliders: 'Window / delay (s)' set to 5.00, 'Long-note threshold (ms)' set to 500, and 'Ensemble timing (± ms)' set to 0. To the right of these sliders is a 'Playback only' checkbox, which is checked. Below the sliders, there is a list of 8 curve lanes on the left, each labeled with a number and a CC number (e.g., 1 · CC 1, 2 · CC 11, etc.). The first lane, '1 · CC 1', is selected. To the right of the list, there are controls for the selected lane: an 'Enabled' checkbox (checked), a 'Target' dropdown menu (set to '1 · Modulation Wheel ...'), a 'Mode' dropdown menu (set to 'Replace'), 'Min' and 'Max' value fields (0 and 127 respectively) with minus and plus buttons, and a 'Glide (ms)' field (30) with minus and plus buttons. Below these controls, there is a 'Curve preset' dropdown menu (set to 'Swell & settle') and a 'Variation %' field (25) with minus and plus buttons. The main area of the interface is a large grid with a blue line representing a curve. The curve starts at the bottom left, rises to a peak in the middle, and then falls to the bottom right. The grid has a light blue background with darker blue lines. At the bottom left of the grid, it says 'note start', and at the bottom right, it says 'note end'. At the very bottom, there is a small text box with instructions: 'Double-click: add point | Alt-click: delete point | Drag between points: bend'.

8 curve lanes

Each lane targets one CC — or the pitch wheel, with full 14-bit bend curves — with its own curve, range and settings.

Replace or Scale

Replace overrides what you play on that CC; Scale multiplies it, keeping your own expression.

Curve editor

Drag points, double-click to add, alt-click to delete, drag between points to bend the segment.

Glide (0–500 ms)

Smooth, legato-aware transitions on the combined output — the CC pre-glides before an overlapping note's attack, and eases back to your value after a curve instead of jumping.

Window & threshold


Set the lookahead and how long a note must be to earn a curve.

Ensemble timing


Spread each note a few ± ms per rack slot, like the players of a section — durations and keyswitches stay exact.

8 factory shapes — plus your own


Swell & settle




Messa di voce




Crescendo




Diminuendo




Attack & sustain



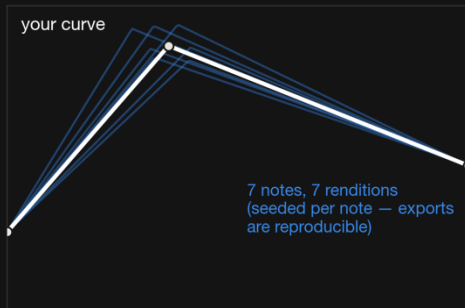

Fade at end



Sforzando



Flat



Variation % — no two notes alike (default 25%)

Per-note random jitter moves the interior curve points slightly in time ($\pm 8\%$) and value ($\pm 10\%$, symmetric around each point), so repeated notes stop sounding copy-pasted — while every note still starts and ends exactly on your curve. Seeded per note: replaying or exporting reproduces the same performance. Save your own edited curves as named presets — available in every project.

Sixteen players, not sixteen copies (new in 0.5.0)

Every slot gets its own rendition

The same settings drive 16 separately-seeded performances: layered libraries stop sounding like one player duplicated.

Curves diverge per slot

The Variation % jitter rolls differently in every slot — each instrument bends the shared curve its own way.

Ensemble timing ($\pm 0\text{--}100$ ms)

Each slot plays every note slightly early or late. Note durations stay exact, keyswitches never move, same-pitch notes never overlap.

Phrase variation (default 10%)

Each slot leans into the melodic contour a little differently — the phrase's reference level stays anchored.


Slot 1 is the reference

Exact timing, exact contour: one rendition always plays precisely what you designed — and single-instrument sessions sound unchanged.

Still fully reproducible

Everything is seeded per note and per slot: replaying or exporting gives the same 16 performances every time.

Phrase humanization — before & after



Applied before the long-note curves — the phrase contour multiplies the final CC output, so it is always audible.

The Phrase tab

HumanizerHost

Expressive MIDI humanization

v0.8.0 · © 2026 Franck Dansaert

Enabled

Phrase

Notes

Instrument

Monitor

Enable phrase humanization

Dynamics CC

1 · Modulation Wheel (coarse) ▾

Vibrato depth CC

Off ▾

100 %

0

Vibrato rate CC

Off ▾

100 %

0

Ignore notes below

C0

Max gap between phrase notes (ms)

100

Silence before a phrase (ms)

1000

Scale per octave above the first note (%)

70

Variation per instrument (%)

10

A phrase is a run of notes closer together than the gap, starting after the silence. The first note sets the reference: the dynamics CC is scaled up for higher notes and down for lower ones (target reached mid-note), and fades to zero on the phrase's last note. Applied before the long-note curves. Notes below the pitch threshold (keyswitches) still play but are invisible to the humanizer: no phrase detection, no long-note curves, not shown in the Monitor.

Phrases are detected automatically

A phrase is a run of notes closer together than Max gap, whose first note follows a real rest (Silence).

Dynamics CC

Pick the CC your library uses for dynamics — usually CC1 (mod wheel) or CC11 (expression).

Scale per octave

How strongly the CC follows the melodic contour: at the default 70%, +1 octave pushes the CC up ×1.7.

Keyswitches stay invisible — everywhere

Notes below the “Ignore notes below” pitch still play, but the whole humanizer ignores them: no long-note curves, no phrase detection, hidden from the Monitor.

Variation per instrument

Each rack slot leans into the contour a little differently (slot 1 plays it exactly).

Vibrato follows the dynamics

Two follower CCs (depth & rate) mirror the dynamics output, each scaled in % — and combine with note curves on the same CC.

The Monitor tab — see what it plays



Before & after, live

One track per monitored CC: the humanized output as a solid curve, what you actually played dashed.

Perfectly aligned

Incoming events are shown at the time their humanized counterpart comes out, so the two curves line up.

Note markers

A mini piano-roll strip shows the note stream — green note-on, red note-off — on the same timeline.

Sliding window & Freeze

View 1–30 s of history (default 5 s) and freeze the display to examine a phrase.

Pick the rendition

A Slot selector switches the view between the 16 per-slot performances.

Zero audio cost

The audio thread feeds a lock-free FIFO only while the tab is open — it never blocks or allocates.

Built for the tutti — multi-core rendering (new in 0.8.0)

One HumanizerHost instance, 16 SWAM instruments + 16 Vienna MIR Pro inserts, 4096 samples @ 48 kHz — an 85 ms budget per block. Time to render one block:

63 ms serial

one instrument after another on the audio thread — 74 % of the budget before anything else runs

10 ms parallel

the slots spread over the performance cores, the audio thread joining in

15 ms worst block

no straggler: no stall from MIR Pro's shared engine, no worker parked on an efficiency core

Performance cores only

The pool has one worker per performance core (12 on an M3 Max, 4 efficiency cores ignored). A slot rendered on an efficiency core took three times longer — and every block waited for it.

Inside the audio workgroup (macOS)

The workers join the audio workgroup of the device your DAW streams on, so the system schedules them with the audio thread instead of as background work. VST3 can't hand it over — HumanizerHost finds the device itself.


Render meters on every slot

Each strip shows its peak render time per block: instrument, then fx. Orange past half the budget, red when the slot alone exceeds it — the red one is the plugin to lighten.

Rack status under the master volume

Peak time of the whole rack against the block's budget, the number of workers and the workgroup they joined. Still bit-identical to serial rendering; the Multi-core toggle turns it off.

Getting started — the Instrument tab



1 · Insert the plugin

Add HumanizerHost as an instrument track in Cubase, Dorico or your DAW.

2 · Scan & load

Manage plugins... → scan for VST3s once. Then Load instrument... and type keywords (e.g. “spitfire strings”) to filter.

3 · Balance the rack

Optionally add an insert FX per slot and set slot / master volumes.

4 · Humanize

Notes tab → enable a lane, pick the CC your library responds to, choose a curve — and play.

5 · Watch the meters

Each strip shows its render time; the line under Master shows the rack against its budget. Orange or red points at the plugin to lighten.

Good to know

Set the threshold ~10% low

DAWs shorten quantized notes slightly — use e.g. 900 ms to catch 1-second notes.

The delay is a feature

It is not reported as latency: the offset is identical on the timeline and in exports. Nudge tracks earlier if needed. All instruments in your template should be hosted in the plugin, so delay will apply to all tracks even those not humanized.

Playback only (default on)

Transport stopped → MIDI passes straight through, no delay: audition patches normally.

Bypass ALL instances

The green/red button above the tabs disables the humanizer in every instance at once — perfect for A/B-ing. Always off at launch.

Notes longer than the window

Still humanized: the curve is stretched to the window and the final value held to the real note-off.

Crash-proof scanning

Each plugin is scanned in a disposable helper process — a crashing plugin cannot take down your DAW.