The inside of the song.

A pre-authorised, documented, non-public corpus of how music is made: 9,463 real DAW sessions with their stems, MIDI and presets. XIM Audio Inc. licenses the archive of We Make Dance Music, online since 2012. Catalogue, measured and documented: ximaudio.com

Music is made in parts, one player and one mixer channel at a time, and a model learns it best the same way, not from finished stereo tracks. Every creator has already authorised AI training, in writing. Nothing here needs a consent process. It is not on the internet. These sessions sit in private download archives bought by producers, never crawled. It is long-form human intention, not snippets: each session is a producer's full decision chain from the first kick to the master. It is multi-modal and aligned: audio, MIDI, preset, mixer state and session structure of the same track. And it is documented per asset, with the creator paid on every licence.

THE LARGEST LICENSED CORPUS OF REAL DAW SESSIONS WITH ALIGNED STEMS, MIDI AND PLUGIN CHAINS

9,463
DAW sessions: Ableton, Logic, FL Studio, Cubase, every one on the live catalogue

98,246
stems in 3,923 songs, printed from the session, never mixed together

287,073
MIDI files, from the packs and from inside the sessions

167,491
synth presets: Serum, Sylenth1, Massive, Spire, Vital, u-he

453,574
labelled loops and one-shots

2,006
acapella and instrumental pairs, each printed from one session

8,153 h
of stem audio, every lane added up. 24-bit WAV, key, tempo, genre set by the producer

3,146
tracks with session, stems and MIDI, all aligned

WHY IT EXISTS NOWHERE ELSE

- **Ground truth.** The stems are the parts the producer printed from the session, not a separation model's output fed back as truth. They add up to the unmastered mix that ships beside them, the unity-gain sum of the lanes, residual recorded per song. No bleed: a lane is printed inside the session, not recorded in a room. A mastered reference on request.
- **Aligned.** Session, stems, MIDI and presets belong to the same track. Many songs exist as a Logic, an Ableton and an FL Studio version of one piece of music.
- **Labelled, lane by lane.** 24 lanes a song at the median, 12 to 45 for eight in ten. Each lane carries one of 53 instrument labels, 71% from the producer's own lane name, 26% from an audio classifier, 4% unlabelled and said so. By lane: drums 31%, keys and synths 18%, other melodic 15%, bass 12%, percussion 9%, vocals 5%, piano, strings, guitar and wind the rest. Peak and activity ratio per lane flag thin lanes. Key on 90% of tonal parts; tempo, genre and DAW on every song. Human-made.
- **Ready to load.** Stems ship in the MoisesDB layout with a data.json per song; the moisesdb loader reads a delivery as is. Checksums, a creator-level train, val and test split that holds across deliveries, DAW versions of one song grouped, a datasheet. Our parser turns every session into one JSON schema: tracks, plugin chains, mixer state, clips, notes, automation.

WHAT IT TRAINS

- Stem separation on parts that were never mixed, paired with their unmastered mix. The public ground-truth sets are research-only and small: MUSDB18 150 songs, MoisesDB 240. The record-setting models trained on 500 to 1,500 private songs nobody can license. We hold 3,923 songs of modern music made in the DAW: deep, progressive and electro house, trance, ambient, breaks and techno, the styles public models score worst on, with pop, hip-hop and rock beside them.
- Generative models that learn structure, not only sound.
- DAW assistants that arrange, mix and route like a producer.
- Mixing and mastering from channel settings and master pairs.
- Transcription and sound design from MIDI and presets beside their own audio.

YOUR DILIGENCE, ANSWERED

Rights per asset	One record per file: owner, product, licence basis, payment status, withdrawal status. Producible on request, not a catalogue-level promise.
Training named	Our creator terms name machine learning training as a permitted partner licence.
Chain of acquisition	Creator uploads to our marketplace under those terms; we license to you. Two steps, both documented.
Creators paid	A share of every licence goes to the creators whose work is in it, pro rata by audio minutes.
Personal data	None in the audio or sessions. No song title, artist, producer or product name in any delivery: songs and creators are pseudonymous IDs, files are numbered lanes with an instrument label.
AI-made content	None. Human producers, human performances.
Documentation	A datasheet ships with every set: sources, dates, processing, de-duplication, exclusion ledger, known limits.
Licence	Non-exclusive. Perpetual, irrevocable training rights. Models and outputs are yours. No redistribution of the recordings. Signed twice in 2026.
Delivery	Private S3 prefix, scoped key per deal, access logs, 30-day window. MoisesDB layout, manifest, checksums, split and datasheet in every delivery. FLAC on request.

Ask for the evaluation set. Four songs: the session, the stems, their unmastered mix, the MIDI and the presets, with the labels, the per-asset records and the index of the whole catalogue. One email, access the same week.
Mikas, founder of XIM Audio Inc. · mikas@ximaudio.com · ximaudio.com/data

Figures counted 30 October 2026 from the live store, as listed per set at ximaudio.com/data. Stem hours are every lane's length added up from the file headers. MIDI, presets and loops are counted file by file inside live products, sessions included.