Dylan Cantor

Brooklyn, NY · (610) 716-6913 · cantor.dylan@gmail.com
dylancantor.com · linkedin.com/in/dylan-cantor-909b77352 · github.com/djc8582

Education

New York University Expected May 2027
Double major in Music Technology (Steinhardt) and Computer Science (CAS); Minor, Mathematics *GPA 3.69*

NYU Paris — IRCAM program Spring 2026
Spatial audio at IRCAM, one day weekly: ambisonics, directivity, binaural rendering *Paris, France*

Relevant coursework. Digital Signal Processing; Operating Systems; Artificial Intelligence; Basic Algorithms; Linear Algebra; Probability & Statistics; Analog & Digital Electronics with labs; Analog & Digital Recording; Electronic & Computer Music Synthesis; MIDI Technology II (Max/MSP).

Research Experience

Research Team Member — Georgia Tech Center for Music Technology May 2026 – Present
Prof. Gil Weinberg's research group *Atlanta, GA (remote)*

- Sole engineer on the iOS client for HearMe, a shipped application that generates a full arrangement from a musician's own playing, returning separable stems, chords, and time-aligned lyrics.
- Built a Web-Audio-style playback engine in React Native with sample-accurate scheduling for the piano and drum interfaces, and offline MIDI rendering of generated results.
- Implemented the full input path — recording, upload, keyboard capture, lyric entry — into the generation service, with a background watcher that survives app suspension.
- Built the result surface: one transport synchronized across stems, chords, lyrics, and play-along mode, with per-project version history. 200+ commits. Joined for a 150-hour summer term and have remained on the team since, meeting weekly with the group and separately with Prof. Weinberg.

Professional Experience

Cofounder and Sole Engineer — Clubs February 2026 – Present
Creative-collaboration platform, live on the iOS App Store *Brooklyn, NY*

- Built and shipped a platform where projects are posted as cards, discovered by swiping, and applied to; accepted collaborators work in real-time channels. 1,976 users and 311 active public projects from 213 creators across 49 schools, with 24,829 swipes, 3,596 applications and 2,262 messages in the last 30 days.
- Designed the AI pipelines: free-text self-descriptions parsed into interests and skills, content moderation, automatic project tagging, semantic search, and an email path that publishes a forwarded message as a card.
- Built a personalized discovery feed on a sawtooth-staircase ranking algorithm, batched and preloaded in the background so the first card paints instantly, plus real-time group and direct messaging with unread tracking, push notifications and universal deep links.
- Sole engineer across a multi-package monorepo spanning the iOS client, backend, admin console and marketing site; 2,300+ commits. Covered in the *Washington Square News*. Cofounded with Nick Benak, who leads design and marketing.

Projects

Jamarama — real-time multiplayer jam system; Winner, Audiotool Challenge June 2026

- Led the audio architecture at the Google DeepMind Magenta hackathon, with Sai Nellutla. Two to four players each control one dimension of a single fused track from their phones, and one player's chord change re-roots the whole band on the next downbeat.
- Placed the master clock in the host browser, fanning beats over WebSocket to phones acting as thin controllers; the lead keyboard fires a local tone in under 20 ms so playing never waits on the network.
- Measured the generative model for tempo stability, found it drifts within roughly 30 seconds with no tempo input, and removed it from the clock entirely — a deterministic clock-locked band wrapped in an arhythmic texture layer that floats off the grid, where drift is inaudible.
- Ran that layer as Magenta RealTime 2 on-device via MLX at $\sim$2.37$\times$ real-time, conditioned on blended player taste, live chord, and crowd energy. Tone.js, WebSocket, Node, Python, MLX.

"I Feel Like" — second-order ambisonic production and binaural mix Spring 2026

- Produced a multitrack piece entirely in the ambisonic domain at IRCAM: about forty sources individually encoded into a second-order field and positioned by azimuth and elevation rather than panned.
- Ran room encoding, directivity shaping and a feedback-delay-network reverb inside the ambisonic domain rather than on a stereo bus, then decoded the same scene two ways — loudspeaker array by all-round ambisonic decoding, headphones by binaural decoding.
- Recorded all source material: field recordings on omnidirectional capsules, instruments from the Baschet brothers' *Structures Sonores* sound sculptures, and tetrahedral-microphone material converted A-format to B-format before encoding. Reaper, IEM Plug-in Suite, Sennheiser AMBEO.

Creative and Performance Work

Growing — character-based act with Miller Fishback, on Cooking Pot 2022 – Present

- Each alias is a character with its own register, letting the project move between rap, avant-garde jazz-rock and electronic music. Seven names on Spotify and Bandcamp, including *rat's nest* (2024).
- Roughly 990,000 Spotify streams across the catalogue, including a 650,000-stream collaboration and released work with Open Mike Eagle, Quelle Chris and Denmark Vessey.

Close Quarters — band with Nick Benak, on Cooking Pot Present

- *EDUCATION*, a 49-minute continuous psychedelic rock concept album about going through school at NYU. Written, recorded and produced together.

Technical Skills

Audio and DSP: Max/MSP, RNBO, Web Audio API, JUCE, ambisonics and binaural rendering, real-time and streaming audio, sample-accurate scheduling
Machine learning: generative audio, music information retrieval, on-device inference (MLX)
Languages and frameworks: Python, TypeScript, Swift, C++; React Native, Node, Supabase
Tools: Git, Reaper, Ableton Live, Logic Pro, IEM Plug-in Suite