

# YAHIR HERNANDEZ

yahirh@mit.edu | yahirh.com | linkedin.com/in/yahirh | 773-387-7267

## EDUCATION

---

**Massachusetts Institute of Technology** – Cambridge, MA

*Master of Science in Music Technology and Computation*

May 2027

*Bachelor of Science in Computer Science & Engineering and in Music (GPA: 4.6)*

May 2026

**Music Tech Courses:** Fundamentals of Music Processing, Interactive Music Systems, Intro To Music Tech

**AI & ML Courses:** Deep Learning, Machine Learning for Signal Processing, Advances in Computer Vision, Intro to ML

**CS Courses:** Data Structures & Algorithms, Software Construction, Software Design, Computer Systems Engineering

## SKILLS

---

**Programming Languages:** Python, JavaScript, TypeScript, Java, Lua, C#, HTML, CSS, C, BluespecSV

**Libraries & Tools:** Librosa, Torchaudio, Tone.js, NexusUI, pretty\_midi, PyTorch, NumPy, SciPy, Matplotlib, Git

## WORK EXPERIENCE

---

**Time-line.io** – Remote

Aug 2025 – Present

*Machine Learning Engineer*

- Spearheaded exploratory data analysis on over 400 tracks, quantifying section-length patterns, drop-detections, pre/fake-drop rates to create reliable drop estimation when model analysis is unavailable.
- Co-developed end-to-end pipeline for the beat-aligned drop-detection model. Refined per-beat features & multi-scale windows, and designed expanded beat-level features.
- Developed beatgrid & tempo algorithm on single-tempo consistent tracks with various tempo analysis approaches..

**Autonomy and Embedded Robotics Accelerated (AERA) Lab** – MIT - Cambridge, MA

Sep 2023 - Present

*Machine Learning Researcher*

- Designed a VR framework in Unity that leverages human-gaze as an attention-aware modality in collaborative audio event localization for human-robot teaming.
- Developed sound event & localization (SELD) model using FOA+IV features with a CRNN + BiGRU architecture.

**Studio Wedge** – Remote - Chicago, IL

Jun 2025 – Feb 2026

*Software Engineer, Lead Game Programmer*

- Led the initiative to outline the programming implementation timeline and integration of LLMs in the studio's game development design framework for streamlined development and personalized debugging in scripting.
- Developed a first-pass model of an enemy path-finding ML model to target players, achieving a ~20% performance increase versus randomized search during testing. Prototyped ideas for a reactive ML music model during gameplay.
- Implemented core gameplay mechanics in Unity: state machines, collision systems, and modular controller scripts

**Collaborative Music Lab** – MIT - Cambridge, MA

Oct 2025 – Dec 2025

*Music Tech Researcher*

- Enabled collaborative music-making across varying musical ability levels by engineering modular, real-time input features in the Creativitas live coding environment, supporting 3 interaction modes: MIDI, sensor-based, & touch.
- Implemented reusable musical prototypes; documented components for future research and public live deployment.
- Corroborated hypotheses of viable collaborative music systems by testing with a cohort of research assistants, incorporating feedback and criticism while researching current practices in collaborative music experiences.

## PROJECTS

---

**Generative Music Diffusion Model** - Project

Jan 2026

- Fine-tuned Stable Audio Open 1.0 latent diffusion model on 500+ track dataset to generate coherent corrido-style instrumentals and tackle issues of cultural representation in generative music models.
- Optimized deep learning training across RTX 4060 (8GB) and Google Colab A100 GPUs with hyperparameter tuning; scaled dataset 25x (20 to 500 tracks) to reduce overfitting and improve audio quality for musical outputs

**Resurfacr** - Project

Oct 2025

- Built full-stack web app helping Spotify users with 1000+ track libraries rediscover forgotten songs through staleness-based scoring and swipe interface.
- Implemented OAuth 2.0 PKCE & modular backend architecture handling Spotify API limits & deprecated features

**Visually Indicated Music** - Paper

May 2025

- Designed a CNN + LSTM pipeline with MIDI-auxiliary stream to synthesize instrument audio from silent video, tested on URMP dataset; cut initial log-mel spectrogram error by >50%.
- Built an offline preprocessing + sliding-window inference stack that made training 12× faster on NVIDIA RTX 4060.