# Mark Zhou

yuzhenzh@andrew.cmu.edu | www.linkedin.com/in/mark-zhou123 | https://github.com/zhoumark123

## Education

**Carnegie Mellon University** *Pittsburgh, PA* May 2023 - Expected May 2027

Bachelor — Computer Science & Music Technology, *QPA*: 3.81/4.00

*Selected Coursework:*

Distributed Systems (current), Introduction to Machine Learning (current), Introduction to Deep Learning, Probability and Computing, Parallel and Sequential Data Structures and Algorithms, Great Theoretical Ideas in Computer Science

## Work Experience

**Neutone** *Tokyo, Japan* May 2025 – Aug 2025

AI Sound Design and Research Intern

- Trained, converted, and optimized open-source generative models for C++ deployment via TorchScript, enabling integration into Neutone Gen, a plugin showcased internationally at Sónar Festival in Barcelona.
- Architected and built an AI-powered audio plugin in C++ with JUCE, implementing buffers, multi-threading, and synchronization to support real-time audio streaming using an encoder-decoder model.
- Contributed to product design discussions, influencing plugin architecture and guiding integration of new AI-driven features.

**Team Crescendo (Video Game Studio)** *Pittsburgh, PA* Fall 2023 – Present

Audio Programmer, Composer, Sound Designer

- Developed game audio systems in Unity with FMOD, including dynamic BGM switching (battle themes, event-specific music) and randomized soundtrack initiation.
- Coordinated with designers and developers, created and integrated original audio assets by composing music and recording/editing sound effects for in-game regions, dialogues, and items.
- Gathered and incorporated player feedback from playtesting sessions to iteratively improve audio design and enhance user experience

## Project Experience

**Language Models** *(Course Project)* Apr 2025 – May 2025

- Implemented core transformer components (multi-head attention, positional encoding, masking, feedforward layers) from scratch in NumPy.
- Designed and trained transformer models with PyTorch and WandB, including decoder-only (language modeling) and encoder-decoder (ASR) architectures.

**ARCO** Feb 2024 – Feb 2025

- Collaborated with Prof. Roger Dannenberg on ARCO, a sound-synthesis engine for embedded and real-time audio.
- Implemented polymorphic unit generators in C++ with chord detection, spectral centroid, and spectral rolloff.
- Improved a public chord detector's efficiency and added a confidence-calculation algorithm.
- Added real-time timbre-transfer functionality via network communication between C++ and Python processes.

**AI-Generated Music Detector** Aug 2024 – Dec 2024

- Built one of the first open-weight AI-generated music detection models.
- Collected 250+ hours of human-created and AI-generated music across 10 genres for training/validation.
- Architected and trained a CNN with PyTorch; achieved 99.0% validation accuracy.

## Languages and Technologies

**Languages:** Python, C++, C, Go, Standard ML, Assembly, Java
**Frameworks/Tools:** PyTorch, NumPy, TensorFlow, Git, JUCE, FMOD, Unity
**AI Tools:** Cursor, Claude Code, Codex