

Kiran Bhat

(510) 362-1942 | kiranvbhat02@gmail.com | kiranvbhat.com | Long Island City, NY

EDUCATION

Stanford University

Stanford, CA

MA, Music, Science, and Technology; BS, Computer Science; GPA: 3.84

Jun 2026

- **Relevant Coursework:** Audio Signal Processing | Music and AI | Psychoacoustics | Sound Recording Technology | Machine Learning | Computer Vision | NLP | Algorithms | Cryptography | Networking | Operating Systems

TECHNICAL SKILLS

- **Languages:** Python | C++ | C | C# | Julia | Java | JavaScript | HTML | CSS | ChucK | LaTeX | Bash
- **Libraries & Tools:** PyTorch | Git | GCP | AWS | Kubernetes | Claude | Codex | Cursor | Logic Pro X | Pro Tools

EXPERIENCE

SoundPatrol

Dec 2024 - Oct 2026

Audio ML Engineer (Part-time)

Remote

- Prototyped novel distortion-free watermarking for genAI music, with 100% detection in high-entropy generations
- Aligned 1700+ hours of MIDI-audio data for fine-tuning transcription model, filtered out 95% of mismatched pairs
- Created new scanners for YouTube, TikTok, and SoundCloud, increasing relevant dashboard results by over 10x
- Deployed scalable audio tagging model for music/non-music separation with > 99% accuracy

ChucK Development Team

Mar 2024 - Jun 2026

Researcher

CCRMA, Stanford CA

- Developed new features for ChucK, a programming language for real-time sound synthesis and music creation
- Built SMucK, a symbolic music framework for ChucK, adopted by 25+ students in Stanford MUSIC 220B/256A

Wati

Jun 2023 - Aug 2023

Software Engineering Intern

Hong Kong

- Trained ML models on curated contact datasets to predict customer conversions, achieved 92% test accuracy
- Built data pipeline in Python to automatically query YouTrack API and upload ticket information to BigQuery

Stanford University School of Engineering

Mar 2022 - Dec 2022

Computer Science Section Leader (TA)

Stanford, CA

- Instructed and graded 30+ computer science students in Python and C++ in the CS106 program

RESEARCH

Robust Neural Audio Fingerprinting using Music Foundation Models

Aug 2025 - Oct 2025

- Fine-tuned music foundation models (MuQ, MERT) with a nonlinear projection head using contrastive learning
- Achieved 88.2% track-level hit rate and 95.5% segment-level F1, vs. 70.3% and 80.8% for prior SOTA (BEATs)
- Beat Shazam-like Dejavu by 39 points (49.2% hit rate); co-first author, NeurIPS 2025 AI for Music Workshop

SMucK: Symbolic Music in ChucK

Jan 2025 - Jun 2025

- Designed framework for symbolic music in the ChucK programming language; accepted to NIME 2025

PROJECTS

Symbolic and Audio-based Music Identification

Mar 2025 - Jun 2025

- Leveraged DNA strand matching technique to identify songs from symbolic queries, matching 20/20 in a small test
- Searched 2,000 noisy phone recordings with MuQ embeddings, matching all 13 queries that had a true match

Spatialized Doppler Effect Simulator

Dec 2024 - Jan 2025

- Created a 3D simulation of the Doppler effect with custom-built spatialized audio in ChucK
- Implemented spatial sound attenuation, stereo panning, Doppler shift, and basic physics engine

Audio Inpainting with Diffusion

Sep 2023 - Dec 2023

- Utilized pretrained, unconditional audio-generating diffusion model to restore masked speech waveforms
- Repurposed an image inpainting algorithm for audio

Ghostwriter

Feb 2023 - Mar 2023

- Engineered several constrained text generation models to write lyrics for a provided sequence of music notes
- Won best project award in CS224N: Natural Language Processing, the most enrolled CS course at Stanford