Jingyue Huang

yuer867.github.io | jih150@ucsd.edu | Yuer867

Research Directions: Music Representation Learning · Controllable Music Generation · Music Understanding

Education

University of California San Diego — *California, USA*
· *Ph.D. in Computer Science*, advised by Professor Julian McAuley — Sept 2024 - Present

New York University — *New York, USA*
· *M.S. in Data Science* — Sept 2022 - May 2024
· Courses: Deep Learning for Media, Digital Signal Processing, Natural Language Processing

Fudan University — *Shanghai, China*
· *B.S. in Data Science*, GPA 3.7/4.0, Rank: 10/84 — Sept 2018 - June 2022
· Courses: Machine Learning, Fusion of Music and Computer Science, Database System

Experience

Yamaha Corporation — *Hamamatsu, Japan*
· Research Intern | Research mentors: Akira Maezawa and Yuta Kusaka — June 2026 - Sept 2026
· Developed MuSTream, a streaming representation for real-time music understanding. (In submission for ICASSP 2027)

Smule, Inc. — *California, USA*
· Research Intern | Research mentor: Yongyi Zang — June 2025 - Sept 2025
· Developed StylePitcher, a style-following pitch generation framework for expressive singing. (ICASSP 2026 & Best workshop paper in NeurIPS 2026)

Music and AI Lab, National Taiwan University — *Taiwan*
· Research Assistant | Research advisor: Professor Yi-Hsuan Yang — Mar 2023 - Aug 2024
· Developed EMO-Disentanger, an emotion-conditioned music generation framework (ISMIR 2024).

Publications

MuSTream: A Streaming Music Representation for Online Music Understanding
Jingyue Huang, Yuta Kusaka, Julian McAuley, Akira Maezawa
International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2027 (in submission)

ProSing: Zero-Shot Fine-Grained Singing Technique Conversion
Xinhao Guan, Jingyue Huang, Kangrui Chen, Julian McAuley, Zhong-Qiu Wang
International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2027 (in submission)

StylePitcher: Generating Style-Following and Expressive Pitch Curves for Versatile Singing Tasks
Jingyue Huang, Qihui Yang, Fei-Yueh Chen, Julian McAuley, Randal Leistikow, Perry R. Cook, Yongyi Zang
International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2026
*NeurIPS Workshop on Artificial Intelligence for Music, 2025 (**Best Paper Award**)*

MuseTok: Symbolic Music Tokenization for Generation and Semantic Understanding
Jingyue Huang, Zachary Novack, Phillip Long, Yupeng Hou, Ke Chen, Taylor Berg-Kirkpatrick, Julian McAuley
International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2026

REGen: Multimodal Retrieval-Embedded Generation for Long-to-Short Video Editing
Weihan Xu, Yimeng Ma, Jingyue Huang, Yang Li, Wenye Ma, Taylor Berg-Kirkpatrick, Julian McAuley, Paul Pu Liang, Hao-Wen Dong
Conference on Neural Information Processing Systems (NeurIPS), 2025

Emotion-Driven Piano Music Generation via Two-stage Disentanglement and Functional Representation
Jingyue Huang, Ke Chen, Yi-Hsuan Yang
International Society for Music Information Retrieval Conference (ISMIR), 2024

Emotion-Driven Melody Harmonization via Melodic Variation and Functional Representation
Jingyue Huang, Yi-Hsuan Yang
arXiv preprint arXiv:2407.20176, 2024

Weakly Supervised Learning for Textbook Question Answering
Jie Ma, Qi Chai, Jingyue Huang, Jun Liu, Yang You, Qinghua Zheng
IEEE Transactions on Image Processing, 2022

Generative Entity Typing with Curriculum Learning
Siyu Yuan, Deqing Yang, Jiaqing Liang, Zhixu Li, Jinxi Liu, Jingyue Huang, Yanghua Xiao
Conference on Empirical Methods in Natural Language Processing (EMNLP), 2022

Large-scale Multi-granular Concept Extraction Based on Machine Reading Comprehension
Siyu Yuan, Deqing Yang, Jiaqing Liang, Jilun Sun, Jingyue Huang, Kaiyan Cao, Yanghua Xiao, Rui Xie
International Semantic Web Conference (ISWC), 2021

Projects

MuSTream (GitHub, Demo)
- *Mentors: Akira Maezawa, Yuta Kusaka | Yamaha Corporation*
- Proposed a streaming music representation for general-purpose online music understanding by distilling the pretrained non-causal MERT model through a two-stage framework.
- Developed a real-time inference interface supporting on-device chord, beat, and key estimation on edge hardware.

StylePitcher (Demo)
- *Mentor: Yongyi Zang | Smule Labs*
- Proposed a pitch curve generator that learns singer style from reference audio, supporting diverse singing tasks including pitch correction, singing voice synthesis and singing voice conversion.
- Introduced a flow matching architecture for curve generation and an inpainting mechanism for flexible task adaptation.

MuseTok (GitHub, Demo)
- *Advisor: Julian McAuley, Professor | University of California San Diego*
- Proposed a tokenization framework of symbolic music, applicable to both generation and semantic understanding tasks.
- Analyzed how music tokens learn underlying musical concepts, such as key, interval, time signature, and texture.

EMO-Disentanger (GitHub, Demo)
- *Advisor: Yi-Hsuan Yang, Professor | National Taiwan University*
- Developed a two-stage framework for emotion disentanglement in piano performance generation, modeling valence through lead sheet composition and arousal via performance-level attributes.
- Proposed a novel functional representation for symbolic music, encoding both melody and chords with Roman numerals relative to musical keys, to consider the interactions among notes, chords and tonalities.

Scholarships and Leaderships

- **Outstanding Student Scholarship : 2021 - 2022, 2020 - 2021, 2019 - 2020**
- Vice President of Technology & Education: Connecting Cultures (TECC), Fudan University

Skills

- **Programming**: Python, SQL, C, R, and Matlab
- **Framework**: PyTorch, PyTorch Lightning
- **Language**: Native in Chinese; Fluent in English
- **Instrument**: Guitar, Guzheng