

Lejun Min

Researcher, Artist

lejun@ccrma.stanford.edu | [Webpage](#) | [LinkedIn](#)

EDUCATION

Center for Computer Research in Music and Acoustics, Stanford University

Sept. 2024 – Present

Master of Art in Music, Science, and Technology (Fellowship)

California, United States

- GPA: 4.0.
- Advisor: Prof. Julius O. Smith, Prof. Takako Fujioka.

ACM Honor Class, Shanghai Jiao Tong University

Sept. 2019 – June 2023

Bachelor of Engineering in Computer Science (Fellowship)

Shanghai, China

- An elite CS program for **top 5%** students.
- GPA: 88.5/100 (ranking: 6/30).

PUBLICATIONS

X. Qu, Y. Bai, Y. Ma, Z. Zhou, K. Lo, J. Liu, R. Yuan, **L. Min**, X. Liu, T. Zhang, X. Du, S. Guo, Y. Liang, Y. Li, S. Wu, J. Zhou, T. Zheng, Z. Ma, F. Han, W. Xue, G. Xia, E. Benetos, X. Yue, C. Lin, X. Tan, S. Huang, W. Chen, J. Fu, G. Zhang, “MuPT: A Generative Symbolic Music Pretrained Transformer”, in *Proc. 13th International Conference on Learning Representations (ICLR 2025)*. [[arXiv](#)] [[OpenReview](#)] [[Demo](#)]

Z. Wang, **L. Min**, G. Xia, “Whole-song Hierarchical Generation of Symbolic Music Using Cascaded Diffusion Models”, **Spotlight (top 5%)** in *Proc. 12th International Conference on Learning Representations (ICLR 2024)*, Vienna, May 2024. [[arXiv](#)] [[OpenReview](#)] [[Demo](#)]

L. Min, J. Jiang, G. Xia, J. Zhao, “Polyffusion: A Diffusion Model for Polyphonic Score Generation with Internal and External Controls”, in *Proc. 24th International Society for Music Information Retrieval Conference (ISMIR 2023)*, Milan, November 2023. [[arXiv](#)] [[Poster](#)] [[Demo](#)]

RESEARCH EXPERIENCE

Automatic Mixing with Audio Representation Learning and Generation

June 2025 - Sept. 2025

Research Intern at Sony CSL Paris

Paris, France

- Achieved end-to-end audio mixing with a two-stage paradigm: mix-invariant representation learning, and latent audio generation. This is one of the pioneer works on this topic.
- Experimented adding different schemes of self-supervision to the representation learning stage.
- Advisor: Dr. Stefan Lattner.

Hierarchical Generation and Performance Rendering of Symbolic Music

Sept. 2023 – Feb. 2024

Research Assistant at Music X Lab, MBZUAI

Abu Dhabi, United Arab Emirates

- Designed and implemented comprehensive experiments for the hierarchical generation of symbolic music, with a cascaded diffusion model as backend.
- Experimented on performance rendering for symbolic music using Transformer architecture.
- Advisor: Prof. Gus Xia.

Controllable Symbolic Music Generation with Diffusion Models

June 2022 – Dec. 2022

Research Assistant at Music X Lab, MBZUAI

Abu Dhabi, United Arab Emirates

- Achieved state-of-the-art polyphonic music generation using diffusion models.
- Devised two control paradigms for music generation in the diffusion model framework: internal control via masked generation, and external control via cross-attention mechanism.
- Advisor: Prof. Gus Xia.

Deep Learning on Piano Reduction and Orchestration

Jan. 2022 – May 2022

Researcher at Music X Lab, New York University, Shanghai

Shanghai, China

- Projected piano and orchestral scores to a joint latent space with variational autoencoders.
- Applied contrastive learning on the latent space with end-to-end autoencoder training.
- Advisor: Prof. Gus Xia.

LANGUAGE PROFICIENCY

Mandarin Chinese (native), English (fluent), French (beginner)

TOEFL: 112 (Reading 30, Listening 30, Speaking 24, Writing 28)

GRE: Verbal 162, Quantitative 170, Writing 4.0

PROGRAMMING PROJECTS

Computer Graphics

Gigantic Splight (Python)

June 2022

An interactive 3D fluids simulation based on Taichi framework.

Scotty3D (C++)

Mar. 2022

A comprehensive CG project including software rasterization, interactive mesh editing, realistic path tracing, and dynamic animation.

Ray Tracer (Rust)

Aug. 2020

A complete ray tracing engine.

Audio Signal Processing

Simple EQ (C++)

Jan. 2022

A step-by-step JUCE learning project for audio plugin development.

Audiobia (Python & Tensorflow)

May 2021

Audio classification using Google's EfficientNet and Harmonic Percussive Source Separation (HPSS).

Compiler, Computer Architecture & System

Mx Compiler (Java)

May 2021

A completely hand-made compiler for a toy language (Java subset) that surpasses -O1 optimization.

RISC-V CPU (Verilog)

Dec. 2020

An emulated 5-pipelined RISC-V CPU with real-world FPGA implementation.

Python Interpreter (C++)

Feb. 2020

A Python language interpreter.

Software Development

Train Ticket System (C++)

June 2020

A cooperated project including backend coding, B+ Tree data structure implementation and frontend website design.

ART PRACTICES

Live Performance & New Media Art

A Chan Conversation

May 2025

A sonic conversation with an ancient Chan Buddhist monk. A Live performance that explores spatialized sound perception with Ambisonics. Performed on CCRMA Open House Concert 2025.

Interface / Narrative Design

Kandinsky Sonified (*Chuck* & *ChuGL*)

Nov. 2024

An interactive audiovisual music sequencer that creates and sonifies Kandinsky-like abstract paintings.

Fireflies (*Chuck* & *ChuGL*)

Oct. 2024

An interactive music therapy journey embodying a firefly. Essentially a sound peeking visualization.

Music

忆久 (**Memories Last Long**)

June 2023

A song and a music video dedicated to the graduates of 2023, Zhiyuan College.

Should Have Known Better (**piano & synth cover**)

Feb. 2023

Piano, synth & singing performance.

晚海 (**Sunset Sea**)

Dec. 2021

A single published under CEM Records, one of the most prestigious electronic music labels in China.

TEACHING

Reinforcement Learning (CS3316)

Spring 2023

Teaching Assistant at SJTU

Shanghai, China

- Designed the final project involving single- or multi-agent learning for simulated hands and legged robot.
- Lecturer: Prof. Weinan Zhang.

Design and Analysis of Algorithms (AI2615)

Spring 2022

Teaching Assistant at SJTU

Shanghai, China

- Lecturer: Prof. Chihao Zhang.

Principle and Practice of Computer Algorithms (CS1952)

Summer 2021

Teaching Assistant at SJTU

Shanghai, China

- Designed a comprehensive ray tracing tutorial written in the Rust language. The repository received 100+ stars on GitHub.
- Supervisor: Prof. Yong Yu.