

Mingjing (MJ) Yi

(646) 319-2375 | my2931@columbia.edu | [linkedin.com/in/mingjing-yi/](https://www.linkedin.com/in/mingjing-yi/) | github.com/ImageYi | [Imageyi.github.io/](https://imageyi.github.io/)

EDUCATION

Columbia University

Sept 2025 - Dec 2026

M.S. in Computer Science, GPA 4.0/4.0

New York, NY

- Coursework: LLM Based GenAI, Deep Learning, Computer Graphics, Natural Language Processing, Continual Learning and Memory Models

Duke Kunshan University (DKU) and Duke University

Sept 2021 - May 2025

B.S. in Computer Science, GPA 3.7/4.0

Kunshan, CN

B.S. in Interdisciplinary Studies; Computer Science (Duke)

Durham, NC

- Honors & Awards: UG Entrance Scholarship (75%), Dean's List (4x)
- Coursework: Artificial Intelligence; Design and Analysis of Algorithm; Elements of Machine Learning; Computer Vision; Speech Recognition

RESEARCH EXPERIENCE

Mobile X Lab at Columbia

Sept 2025 - Present

Research Assistant Co-advised by [Prof. Xia Zhou](#) and [Prof. Changxi Zheng](#)

New York, NY

- Developed RollingPol, a single-shot computational imaging system that utilizes a custom VAE-based reconstruction network to extract physics-based polarization priors, significantly improving the noise suppression and visual fidelity of state-of-the-art deep learning models for dynamic underwater image enhancement. Conditionally accepted to **MobiCom 2026** (RollingPol: Single-Shot Polarization Imaging for Underwater Sensing).
- Currently working on 3D reconstruction for underwater scenes with low light and low visibility conditions.

SeeleAI

May 2025 - Present

Algorithm Research Intern

Shenzhen, CN

- Architected EVA01, a 3D Large Multimodal Model leveraging a novel Mixture-of-Transformers (MoT) architecture to unify 3D object understanding, generation, and editing within a single framework. Contributed to core architectural design and executed end-to-end model training and large-scale experimental validation. Co-first author of the paper conditionally accepted to **ACM Transactions on Graphics (SIGGRAPH Asia 2026)** ([EVA01: Unified Native 3D Understanding and Generation via Mixture-of-Transformers](#)). Following work includes a unified representation of 3D mesh and large-scale dataset for 3D instruct tuning.
- Worked on text-to-motion (T2M) generation task. Implemented SOTA methods and trained model for motion generation. Co-designed the rules and topics for motion capturing data and expanded the dataset (~20k motion) to improve generalization. This work is integrated into a paper conditionally accepted to **ACM MM 2026** (MBench: A Large-scale Benchmark for Human Motion Retrieval and Generation).

SMIIP Lab at DKU

May 2024 - May 2025

Research Assistant advised by [Prof. Ming Li](#)

Kunshan, CN

- Conducted research on Video-to-Audio (V2A) generation, leveraging multimodal models such as CLIP and CLAP. Built diffusion-based V2A model with visual scene detection. For scene detection, our method outperforms current tool by over 50% in accuracy. The overall generation quality is increased by 24% in relevance and fidelity compared to SOTA models. Accepted to **APSIPA ASC 2025**: [Efficient Video to Audio Mapper with Visual Scene Detection](#).

PROFESSIONAL EXPERIENCE

Microsoft

May 2026 - Present

Applied Scientist Intern at Frontier Tuning

Redmond, WA

- Evaluated Frontier Tuning levers on enterprise datasets. Implemented and evaluated test-time compute approaches and different tuning methods from context (GEPA, ACE), orchestration, rubrics to model tuning by RL across real business scenarios. Contributed to designing an enterprise benchmark for agentic AI system.

LEADERSHIP & SERVICE

The Division of Natural and Applied Science at DKU

Jan 2023 - May 2023; Jan 2024 - Dec 2024

Teaching Assistant

Kunshan, CN

- Held office hours for over 50 times, assisting students with course material and assignments for *Intro to Data Science* and *Intro to Programming and Data Structures*. Assisted Graduate TAs with recitation sessions.

ADDITIONAL INFORMATION

Languages: Chinese (native)

Skills: AI/ML: PyTorch, Diffusion models, LLMs, AI-native workflow. Systems: CUDA, Distributed training (DDP, FSDP, SLURM), Git