

Mingjing (MJ) Yi

(646) 319-2375 | my2931@columbia.edu | [linkedin.com/in/mingjing-yi/](https://www.linkedin.com/in/mingjing-yi/) | github.com/1ImageYi

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. 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

- Contributed to EVA01, a 3D Large Multimodal Model with Mix-of-Transformers (MoT) architecture that allows a unified model for 3D object understanding, generation and editing. Mainly contributed to design of the architecture and most of the training and experiments. Submitted to **SIGGRAPH Aisa 2026** as first Co-author (EVA01: Empowering Multimodal LLMs for Context-Aware 3D Generation via Mixture-of-Transformers).
- 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 submitted to **ACM MM 2026** (MBench: A Large-scale Benchmark for Human Motion Retrieval and Generation).

SMIP Lab at DKU May 2024 - Aug 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](#).
- Initiated further research exploring the potential of Diffusion Transformer (DiT) architecture as a generative model and add ControlNet for temporal alignment control.

PROFESSIONAL EXPERIENCE

Microsoft May 2026 – Present

Applied Scientist Intern at Frontier Tuning Redmond, WA

- Built an automated bug-filling skill for AI agents shipped in GitHub Copilot and prototyped in internal agents.
- Benchmarked Copilot Tuning levers on private enterprise datasets. Implemented and evaluated test-time compute approaches and different tuning methods from context, orchestration, rubrics to model tuning by RL across real business scenarios.

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. Systems: CUDA, Distributed training (DDP, SLURM), Git