

# Fangzhou Mu (穆方舟)

Email: [fm2@wisc.edu](mailto:fm2@wisc.edu) | Mobile: +1 (608) 332-9265 | [Personal webpage](#) | [Google scholar](#) | [GitHub](#) | [LinkedIn](#)

## SUMMARY

**Research Interests:** Foundation Models, Computer Vision, Machine Learning

- *Strong knowledge in multimodal foundation models, diffusion models, 3D vision, and video understanding.*
- *Extensive experience in the design, implementation, training, and performance optimization of deep models in PyTorch.*
- *Publication in top-tier vision and ML conferences (CVPR x6, ICLR x2, ICCV x1) and journals (TPAMI x1).*

## EDUCATION

**University of Wisconsin-Madison**

**Madison, WI, USA**

Ph.D. in Computer Sciences | Advisor: Prof. [Yin Li](#)

Sept. 2018 – Dec. 2023

M.S. in Computer Sciences & M.S. in Pharmaceutical Sciences

Sept. 2014 – Jun. 2018

**Awarded University of Wisconsin Distinguished Graduate Fellowship**

**Zhejiang University (浙江大学)**

**Hangzhou, China**

B.S. in Biological Sciences (GPA: 3.86/4.0, Ranking: 1/33)

Sept. 2010 – Jun. 2014

**Awarded National Scholarship (top 3%)**

## WORK EXPERIENCE

**NVIDIA (AI Foundation Models)**

**Santa Clara, CA, USA**

Senior Deep Learning Algorithms Engineer

Jan. 2024 – Present

**Performance optimization of (multimodal) large language models**

- *Build automation of NVIDIA Inference Microservices (NIM) for serving vision-language models (VLMs) in the cloud.*
- *Performance profiling and optimization of LLM and VLM inference engines for >10x higher generation throughput.*

**University of Wisconsin-Madison**

**Madison, WI, USA**

Graduate Research Assistant (Advisor: Prof. [Yin Li](#))

Sept. 2018 – Dec. 2023

**Thread 1: 3D vision with single-photon LiDAR sensors** (with Prof. [Andreas Velten](#) and Prof. [Mohit Gupta](#))

- *Invented NeRF-based deep models for high-speed imaging beyond the line of sight (ICCP/TPAMI 22).*
  - *Developed learned spatiotemporal compression of SPAD\* histograms for 3D imaging (ICCV 23, US patent filed).*
  - *Developed 3D reconstruction method using distributed SPAD sensors and differentiable transient rendering (CVPR 24).*
- (\* SPAD: single-photon avalanche diode, an emerging time-of-flight sensor for 3D imaging with single-photon light sensitivity and picosecond-scale time resolution.)

**Thread 2: Scalable and resource-adaptive video understanding**

- *Proposed scalable Transformer design and efficient training techniques for action detection in long videos (CVPR 24).*
- *Won top prizes in prestigious video action detection challenges (1<sup>st</sup> prize: EPIC-Kitchens 2023, 2<sup>nd</sup> prize: Ego4D 2022 and 2023).*
- *Developed a latency-aware scheduler for resource-adaptive video object detection on mobiles (CVPR 22).*

**Thread 3: Controllable text-to-image generation with Diffusion Models**

- *Proposed a training-free method for the spatial control of Stable Diffusion with any conditions (CVPR 24).*
- *Introduced adapter-based fine-tuning of Stable Diffusion for grounded text-to-image generation. (CVPR 23).*

**Other work:** Studied linearized and multi-task fine-tuning for adapting foundation models (ICLR 20, ICLR 24).

**OPPO US Research Center (Computational Photography Team)**

**Palo Alto, CA, USA**

Research Intern (Manager: Dr. [Yanli Liu](#))

Jun. 2023 – Aug. 2023

**Personalization meets controllability in text-to-image diffusion**

- *Develop methods for fast, one-shot personalization of Stable Diffusion without fine-tuning.*
- *Introduce object pose control and camera viewpoint control to personalized Stable Diffusion.*

**Snap Research (Computational Imaging Team)**

**New York, NY, USA**

Research Intern (Managers: Dr. [Jian Wang](#) and Dr. [Yicheng Wu](#))

May 2022 – Aug. 2022

**Image restoration for face captures in the wild**

- *Developed StyleGAN inversion using spatially varying latent codes for improved facial identity preservation.*
- *Invented an image quality aware diffusion model with state-of-the-art face restoration capability.*

**Snap Research (Computational Imaging Team)**

**Remote**

Research Intern (Manager: Prof. [Shree Nayar](#))

May 2021 – Aug. 2021

**3D Photo Stylization** (CVPR 22 oral, mentored by Dr. [Jian Wang](#) and Dr. [Yicheng Wu](#))

- *Invented a method for novel view synthesis from a single image in an artistic style for creative 3D photo browsing.*
- *Proposed a point cloud stylization approach to facilitate multi-view consistency of stylized images.*
- *Implemented custom CUDA ops in PyTorch for differentiable 3D point cloud rasterization.*

**University of Wisconsin-Madison / Morgridge Institute of Research**

**Madison, WI, USA**

**ML4BIO - Machine learning literacy for biologists**

- Developed and delivered an [AI4Science workshop](#) to a broad audience (ISMB 22, GLBIO 19 and [Software Carpentries](#)).
- Built [open-source software with GUI](#) to support interactive learning of [sklearn ML models](#) (`pip install ml4bio`).

**PUBLICATION**

- [15] [Fangzhou Mu](#)<sup>\*</sup>, Carter Sifferman<sup>\*</sup> (equal contribution), Sacha Jungerman, Yiquan Li, Zhiyue Han, Michael Gleicher, Mohit Gupta, Yin Li. **Towards 3D Vision with Low-Cost Single-Photon Cameras**. *CVPR 2024*
- [14] [Fangzhou Mu](#)<sup>\*</sup>, Sicheng Mo<sup>\*</sup> (equal contribution), Yin Li. **SnAG: Scalable and Accurate Video Grounding**. *CVPR 2024*
- [13] Sicheng Mo<sup>\*</sup>, [Fangzhou Mu](#)<sup>\*</sup> (equal contribution), Kuan Heng Lin, Yanli Liu, Bochen Guan, Yin Li, Bolei Zhou. **FreeControl: Training-Free Spatial Control of Any Text-to-Image Diffusion Model with Any Condition**. *CVPR 2024*.
- [12] Zhuoyan Xu, Zhenmei Shi, Junyi Wei, [Fangzhou Mu](#), Yin Li, Yingyu Liang. **Towards Few-Shot Adaptation of Foundation Models via Multitask Finetuning**. *ICLR 2024*
- [11] Yimeng Dou, [Fangzhou Mu](#), Yin Li, Tomy Varghese. **Sensorless End-to-End Freehand Ultrasound with Physics-Inspired Network**. *IUS 2023*
- [10] Felipe Gutierrez-Barragan, [Fangzhou Mu](#), Andrei Ardelean, Atul Ingle, Claudio Bruschini, Edoardo Charbon, Mohit Gupta, Yin Li, Andreas Velten. **Learned Compressive Representations for Single-Photon 3D Imaging**. *ICCV 2023*
- [9] Yuheng Li, Haotian Liu, Qingyang Wu, [Fangzhou Mu](#), Jianwei Yang, Jianfeng Gao, Chunyuan Li, Yong Jae Lee. **GLIGEN: Open-set Grounded Text-to-image Generation**. *CVPR 2023*
- [8] [Fangzhou Mu](#), Sicheng Mo, Jiayong Peng, Xiaochun Liu, Ji Hyun Nam, Siddeshwar Raghavan, Andreas Velten, Yin Li. **Physics to the Rescue: Deep Non-line-of-sight Reconstruction for High-speed Imaging**. *ICCP/TPAMI 2022*
- [7] Chris S Magnano, [Fangzhou Mu](#), Rosemary S Russ, Milica Cvetkovic, Debora Treu, Anthony Gitter. **An Approachable, Flexible, and Practical Machine Learning Workshop for Biologists**. *ISMB/Bioinformatics 2022*
- [6] [Fangzhou Mu](#), Jian Wang<sup>†</sup>, Yicheng Wu<sup>†</sup>, Yin Li<sup>†</sup> (co-corresponding authors). **3D Photo Stylization: Learning to Generate Stylized Novel Views from a Single Image**. *CVPR 2022 (oral presentation)*
- [5] Ran Xu, [Fangzhou Mu](#), Jayoung Lee, Preeti Mukherjee, Somali Chaterji, Saurabh Bagchi, Yin Li. **SmartAdapt: Multi-branch Object Detection Framework for Videos on Mobiles**. *CVPR 2022*
- [4] Yin Li, Runyu L Greene, [Fangzhou Mu](#), Yu Hen Hu, Robert G Radwin. **Towards Video-based Automatic Lifting Load Prediction**. *HFES 2020*
- [3] [Fangzhou Mu](#), Yingyu Liang, Yin Li. **Gradients as Features for Deep Representation Learning**. *ICLR 2020*
- [2] Tingting Liang, Qi Zhao, Shan He, [Fangzhou Mu](#), Wei Deng, Bingnan Han. **Modeling Analysis of Potential Target of Dolastatin 16 by Computational Virtual Screening**. *Chemical and Pharmaceutical Bulletin 2018*
- [1] Robyn Umans, Hannah Henson, [Fangzhou Mu](#), Chaithanyarani Parupalli, Bensheng Ju, Jennifer Peters, Kevin Lanham, Jessica Plavicki, Michael Taylor. **CNS angiogenesis and barrierogenesis occur simultaneously**. *Developmental Biology 2017*

**PATENTS**

- [1] **P210037US02 – Movement Monitoring System**. Robert Radwin, Yin Li, Runyu Greene, [Fangzhou Mu](#), Yu Hen Hu

**AWARDS****Research Competitions**

- 1<sup>st</sup> Prize, EPIC-Kitchens Action Detection 2023 Challenge
- 2<sup>nd</sup> Prize, Ego4D Moment Queries 2023 Challenge
- 2<sup>nd</sup> Prize, Ego4D Moment Queries and Natural Language Queries 2022 Challenge

**Conference Awards**

- CVPR 2023 Outstanding Reviewer Award (among 232 of >7,000 reviewers)
- ICLR 2020 Student Travel Award

**PROFESSIONAL SERVICES**

**Journal Reviewer:** TPAMI, TNNLS, TIP, The Visual Computers (TVCI), Ad Hoc Networks

**Conference Reviewer:** CVPR, ECCV, ICCV, WACV, NeurIPS, ICML, ICLR

**TEACHING AND MENTORSHIP**

**Lead Developer and Instructor:** Machine Learning for Biology (ml4bio) workshop (2018, 2019)

Taught machine learning best practices and biological applications to >20 Ph.D. students and postdocs.

**Graduate Teaching Assistant:** Deep learning for Computer Vision, Operating Systems, Advanced Bioinformatics

**Research Mentorship:** Yiquan Li (M.S. student), Shashank Verma (M.S. student), Siddeshwar Raghavan (M.S. student), Abrar Majeedi (M.S. student), Alex Huang (undergrad), Sicheng Mo (undergrad), Evelin Yin (undergrad)