Yixuan Jia

Ann Arbor, Michigan, USA | jiayx@umich.edu | +1 (734)-846-2853 | Homepage | Google Scholar

Research Interests

Generative Models, Representation Learning, AI for Science, with recent focus on **Video Generation** and **World Models**, which share foundational structure with my work on generative priors and diffusion-based inference.

Education

University of Michigan, PhD candidate in Electrical and Computer Engineering — Jan 2023 – Now

- Advisor: Jeffrey Fessler, Qing Qu

University of Michigan, M.S. in Electrical and Computer Engineering — Sept 2020 – Dec 2022

- GPA: 4.0/4.0
- Coursework: Matrix Methods for SIPML, Optimization Methods for SIPML, DL for Computer Vision, Large Language Models

Tsinghua University, B.S. in Instrument Science and Technology (Primary) — Sept 2014 – Jul 2018

- Core Coursework GPA: 3.8/4.0
- Awards: Scholarship for Future Scholars, 2018, Tsinghua University (Top 3%)

Tsinghua University, B.B.A. (Secondary) — Sept 2015 – Jul 2018

- Core Coursework GPA: 3.8/4.0

Work Experience

ByteDance (Seed), Research Scientist Intern, GenAI for Science – San Jose, CA, USA — May 2026 – Aug 2026

Supervisor: Yuning Shen, Yilai Li, Quanquan Gu

- (In progress) Contributed to the development of large-scale generative foundation models for scientific applications.
- (In progress) Designed and implemented training and evaluation pipelines for diffusion / flow-based generative models.

Woobo Inc., Algorithm & Software Development Engineer Intern – Beijing, China — Sept 2018 – Jun 2019

Supervisor: Feng Tan

- Implemented a speech recognition system using Kaldi to enhance command recognition accuracy.
- Developed a speech synthesis module based on HTS (HMM-based Text-to-Speech) for natural human-robot voice interaction.
- Built and assembled a 3D-printed prototype using SolidWorks for visualization.

Research Experience

University of Michigan, Graduate Research Assistant, Fessler & Qu Labs – Ann Arbor, MI, USA — Jan 2023 – Now

Advisor: Jeffrey Fessler, Qing Qu

I. AI for Science & Video Generation

- Developed a family of diffusion- and flow-based generative frameworks for scientific data assimilation. **FlowDAS (NeurIPS 2025)**: first stochastic-interpolant framework that jointly learns observation-conditioned state-transition dynamics and generative priors for unknown stochastic systems; achieved state-of-the-art on Navier-Stokes and weather nowcasting (SEVIR) benchmarks. **ForcingDAS (ICML 2026 Workshop)**: self-contained data assimilation framework built on diffusion forcing, with a per-frame noise schedule that learns a single joint-trajectory (video) prior spanning the full filtering-to-smoothing spectrum at inference time; evaluated on 2D Navier-Stokes vorticity, SEVIR precipitation nowcasting, and global atmospheric state estimation.
- Co-developed **Data-Forcing Distillation (DFD, ICML 2026 Workshop)** with NVIDIA Research: a distillation method that restores both diversity and fidelity in few-step video generation, enabling efficient sampling without sacrificing generation quality.

II. Representation Learning in Diffusion Models

- Proposed **ICR (ICML 2026)**: the Invariant Contamination Ratio, a label-free metric based on Fisher-style invariance–residual decomposition that identifies optimal feature-extraction noise levels in diffusion models and serves as an early signal of memorization under limited data.
- Co-developed **MCLR (ICML 2026 Workshop)**: a principled alignment objective that maximizes inter-class likelihood-ratios during training, achieving classifier-free-guidance-quality samples under standard sampling, and formally established the equivalence between CFG and alignment-based training.

III. LLM Agents for Scientific Imaging

- Co-developed **Imaging-101 (ICCP 2026 Oral)** with He Sun's group at Peking University: a comprehensive benchmark evaluating

LLM agents on scientific computational imaging tasks, exposing key failure modes including domain-specific scientific knowledge gaps, "scientific debugging atrophy" (defaulting to trial-and-error over diagnostic reasoning), and lack of physics-grounded numerical intuition for units, scaling, and conditioning.

IV. Computational Imaging

- Led 2 journal papers (**SpeRF**, **Y90 SPECT**) and contributed to 2 SNMMI conference abstracts (**At-211 SPECT**, **SAMS**) on deep-learning methods for clinical SPECT imaging, in collaboration with the UMich Nuclear Medicine Department.
- Highlights: **SpeRF (EJNMMI Physics 2025)** achieved up to 4× reduction in clinical Lu-177 SPECT acquisition time via NeRF-style self-supervised projection synthesis (phantom + 17 patient studies); **Y90 SPECT (EJNMMI Physics 2023)** achieved 66% lower NMAE than the Monte-Carlo gold standard via a unified 3-stage CNN for scatter estimation, reconstruction, and dosimetry.

Entrepreneurial Experience

H2OS (h2oswater.com), Co-founder & Lead AI Researcher – Ann Arbor, MI, USA Jan 2026 – Present

Co-founders: Yuhan Li (CEO), Yiting (Sophia) Li (CTO)

What is H2OS? *H2OS is an AI-powered aquaculture startup developing dissolved-oxygen sensing and forecasting systems to help fish farms prevent hypoxia and optimize aeration, feeding, and water-management decisions.*

- Led the design of H2OS's core dissolved-oxygen forecasting model in aquaculture ponds: an LSTM-based time-series model trained via physics-guided transfer learning, using a data-rich estuary dataset to improve forecasting in a data-scarce fish farm setting.
- Incorporated two domain-specific physical constraints into LSTM training: Henry's law (temperature-dependent DO solubility) and the monotonic nighttime DO depletion trend, improving physical consistency and reducing negative transfer across environments.
- Using only the first four nighttime measurements to recursively forecast DO until dawn, the proposed method achieves $> 30\%$ lower RMSE than a standard LSTM, demonstrating the value of combining transfer learning with physics-based regularization.

Publications

* Equal contribution † Corresponding author

First-Author Publications

[6] S. Chen*, J. Ying*, ***Y. Jia****, Y. Gu*, E. Ye, W. Bai, Z. Zeng, S. Ren, B. Gao, Y. Li, T. Zhang, H. Sun†. Imaging-101: Benchmarking LLM Agents for Scientific Computational Imaging. *ICCP 2026* (Oral).

[5] ***Y. Jia***†, S. Chen, Y. Pan, X. Li, L. Shi, C. Jung, H. Yuan, I. Alkhouri, Y. Wu, S. Ravishankar, J. Fessler, Q. Qu. ForcingDAS: Unified and Robust Data Assimilation via Diffusion Forcing. *ICML 2026 Workshop* (Poster).

[4] X. Li*, ***Y. Jia****, Z. Zhang, X. Li, L. Shi, J. Zhou, Z. Zhu, L. Shen, Q. Qu†. Evaluating the Representation Space of Diffusion Models via Self-Supervised Principles. *ICML 2026* (Poster).

[3] S. Chen*, ***Y. Jia****, Q. Qu, H. Sun†, J. Fessler. FlowDAS: A Stochastic Interpolant-Based Framework for Data Assimilation. *NeurIPS 2025* (Poster).

[2] Z. Li*, ***Y. Jia****†, X. Xu, J. Hu, Y. Dewaraja, J. Fessler. Shorter SPECT Scans using Self-Supervised Coordinate Learning to Synthesize Skipped Projection Views. *EJNMMI Physics* | *IMSI Computational Imaging 2024 Workshop* (Oral).

[1] ***Y. Jia***†, Z. Li, A. Akhavanallaf, J. Fessler, Y. Dewaraja. ^{90}Y SPECT Scatter Estimation and Voxel Dosimetry Using a Unified Deep Learning Framework. *EJNMMI Physics* | *SNMMI 2023* (Oral).

Other Publications

[5] S. Chen, S. Liu, ***Y. Jia***, Z. Wang, H. Ling, Q. Qu, J. Gao†. Data-Forcing Distillation: Restoring Diversity and Fidelity in Few-Step Video Generation. *ICML 2026 Workshop* (Poster).

[4] F. Hussain, ***Y. Jia***, Z. Li, Z. Lu, J. Fessler, Y. Dewaraja†. Deep Residual Learning Framework for Scatter Estimation in SPECT Imaging of Alpha Emitters: Application in ^{211}At SPECT. *SNMMI 2026* (Poster).

[3] X. Li†, ***Y. Jia***, X. Li, J. Fessler, R. Wang, Q. Qu. MCLR: Improving Conditional Modeling in Visual Generative Models via Inter-Class Likelihood-Ratio Maximization. *ICML 2026 Workshop* (Poster).

[2] Z. Lu, Z. Li, ***Y. Jia***, G. Chen, M. Roseland, G. Mok, Y. Dewaraja†. Segment Anything Model for SPECT (SAMS): Novel Implementation in SPECT Imaging for Tumor Segmentation. *JNM Supplementary* | *SNMMI 2024* (Oral).

[1] Y. Gou, ***Y. Jia***, P. Wang†, C. Sun. Progress of Inertial Microfluidics in Principle and Application. *Sensors*.

Skills

Research Areas: Diffusion Models, Flow Matching, Generative AI, Foundation Models, Representation Learning, Self-Supervised Learning, Data Assimilation, Inverse Problems, Computational Imaging, Video Generation, World Models

Models & Techniques: Stochastic Interpolants, Score Matching, Diffusion Forcing, Diffusion Transformer (DiT), Classifier-Free Guidance (CFG), Contrastive Learning, Model Distillation, LoRA / Fine-Tuning, Neural Radiance Fields (NeRF), Variational Inference, Bayesian Filtering / Smoothing

ML Frameworks & Libraries: PyTorch, JAX, PyTorch Lightning, Hugging Face Diffusers / Transformers, Weights & Biases

Distributed Training & Infrastructure: Multi-GPU / Multi-Node Training (DDP, FSDP), Docker, Git, Linux, Slurm, HPC Clusters

Programming Languages: Python, C++, C, MATLAB, Julia

Math Foundation: Optimization, Probability & Statistics, Linear Algebra

AI-Assisted Development: Cursor, Claude Code, GitHub Copilot

Teaching

Graduate Student Instructor, University of Michigan — Fall 2022 – Spring 2025

- EECS 559 Optimization Methods for SIPML (Wn23, Wn25)
- EECS 551 Matrix Methods for SIPML (Fa24)
- EECS 453 Principles of Machine Learning (Fa22)
- Biomedical AI (Fa23)
- AI Magic Summer School — K-12 outreach (Su24)

Service

Conference & Journal Reviewer: ICML, CPAL, EJNMMI Physics, QIMS

Organizer, The 2025 Michigan Student Symposium for Interdisciplinary Statistical Sciences (*MSSISS 2025*) — Oct 2024 – Mar 2025