Jiawei (Jimmy) Xu

Email | Linkedin | Website Github

Education

- **University of Maryland, College Park** — *College Park, MD*
  Ph.D. in Applied Mathematics & Statistics and Scientific Computation — September 2025 - Expected May 2029
  Ph.D. Advisor: Prof. Furong Huang

- **New York University** — *New York, NY*
  Honors B.A. in Mathematics — January 2022 - May 2025

Publications and Preprints

* = equal contribution / alphabetical order

[1] LLM Self-Improvement **J. Xu***, M. Liu*, J. Zhang, T. Goldstein, and F. Huang (2026). **β-OPSD: Deriving with Policy Optimization, Training with Self-Distillation**. *In submission to ICLR 2027*.

[2] LLM Reasoning Control **J. Xu**, M. Liu, A. Agrawal, Y. Chen, F. Huang (2026). **Scheduling Thoughts: Learning the Order of Thought in Diffusion Language Models**. *ICML 2026.*

[3] ML Theory Y. Chen*, E. Vanden-Eijnden*, and **J. Xu*** (2025). **Lipschitz-Guided Design of Interpolation Schedules in Generative Models**. *TMLR*

[4] Synthetic Data Generation B. Feuer*, **J. Xu***, N. Cohen, P. Yubeaton, G. Mittal, and C. Hegde (2024). **SELECT: A Large Scale Benchmark of Data Curation Strategies for Image Classification**. *NeurIPS 2024 Datasets and Benchmarks Track.*

Ongoing Projects

- **World Modelling Guided Inference in Multi-turn Agent Tasks** — *UMD CS*
  June 2026 – Now
  - Developing world models for foresight; predicting future agent states and outcomes to guide inference trajectories in multi-turn agentic reasoning tasks.

Selected Research Projects

- **Logit Interpolation in β-OPSD** — *UMD CS*
  Preprint: arXiv:2607.28582 — February 2026 – Augest 2026
  - Developed a generalized framework for on-policy self-distillation via KL-regularized reinforcement learning and introduced logit-interpolation distillation, improving performance by 10% on AIME 2024 and 5% on AIME 2025.

- **Learning the Order of Thought in Diffusion Language Models** — *UMIACS, UMD*
  Accepted to ICML 2026 — September 2025 - January 2026
  - Proposed a dense verifier-free self-aware reward derived from the frozen denoiser's pathwise log-likelihood in masked diffusion models, enabling optimal scheduling of any-order inference for reasoning.
  - Combined the learned scheduling framework with reinforcement learning for reasoning tasks, achieving state-of-the-art results on logic puzzle, math, and coding benchmarks with up to a 20% performance gain.

- **Lipschitz Desgin of Interpolation Schedule in Generative Modeling** — *Courant Institute, NYU*
  In submission to TMLR — September 2024 - Augest 2025
  - Developed a thorough research on scalar interpolation in generative modeling.
  - Derived the analytical optimal lipschitz transport of Gaussian random variables and Gaussian mixture models in high dimensions. Evaluated in numerical detail.
  - Designed adaptive min-lip sampling algorithms and effectively reduced the high-dimensional mode collapse sampling issues with flow matching.

- **SELECT: A Large-scale Benchmark of Data Curation Strategies** — *Tandon, NYU*
  Accepted to NeurIPS 2024 — July 2023 - October 2023
  - Proposed SELECT, a data curation benchmark in the domain of image classification, employing the utility metrics and analytic metrics of a total of 8 evaluations.
  - Generated one of the training datasets with **Clip embeddings** in the ImageNet++, the **largest, most diverse** set of distribution shifts for ImageNet-train to date, including over 6 million samples.

Techincal Skills

- **Programming Languages:** Python, C++, Java, Julia, Matlab
- **Data Science & Machine Learning:** Pytorch, Tensorflow, JAX
- **Other Tools & Technologies:** Github, Linux, Numpy, Scipy, Pandas