

ZIYU LI

ziyul6@illinois.edu | [ziyuli1121.github.io](https://github.com/ziyuli1121)

EDUCATION

- **The Chinese University of Hong Kong (CUHK)** Aug 2026 (Expected) – Jul 2030 (Expected)
Ph.D. in Computer Science and Engineering
- **University of Illinois Urbana-Champaign** Jan 2024 – May 2026
B.S. in Liberal Arts and Sciences, Magna Cum Laude and with High Distinction in Mathematics
Data Optimization Concentration; Minor in Computer Science
 - GPA: 3.98/4.0 (Major GPA: 4.0/4.0)
- **Tongji University** Sep 2021 – Dec 2023
B.Eng. coursework in Mechanical Engineering; transferred to UIUC

RESEARCH INTERESTS

- Multimodal generative modeling, diffusion model, reinforcement learning

RESEARCH EXPERIENCE

- **Research Assistant**, Wave Intelligence Lab @ CUHK Oct 2025 – Present
Advisor: Shengchao Liu
 - Investigating scale-aware generative modeling frameworks that progressively discover, represent, and generate hierarchical structure across coarse-to-fine scales
- **Research Intern**, AGI Lab @ Westlake University May 2025 – Dec 2025
Advisor: Chi Zhang
 - Developed test-time scaling techniques for diffusion and flow matching models
 - Designed efficient ODE solvers using distillation and reinforcement learning for accelerated diffusion sampling
- **Research Assistant**, Carl R. Woese Institute for Genomic Biology Sep 2024 – Dec 2025
Advisor: Huimin Zhao
 - Developed and benchmarked sequence-, structure-, and graph-based models for protein representation learning, including a geometry-aware GNN with electrostatic field features for solubility prediction
 - Prototyped a hierarchical conditional generation framework for *de novo* enzyme design by decomposing biological constraints across cofactor, mechanism, and substrate levels

PUBLICATIONS

- **Parallel Diffusion Solver via Residual Dirichlet Policy Optimization**, *Co-first Author* IEEE TPAMI, 2026
[arXiv:2512.22796](https://arxiv.org/abs/2512.22796) | [Project Page](#)
 - Developed **RDPO** (Residual Dirichlet Policy Optimization), a parameter-efficient RL framework that reformulates diffusion sampling as a stochastic policy in a simplex-constrained space
 - Modeled the **EPD-Solver** updates as a Markov Decision Process (MDP), employing a Residual Dirichlet Policy to adaptively learn optimal intermediate timesteps and gradient mixture weights
 - On Text-to-Image benchmarks, our RDPO-tuned EPD-Solver significantly improves human preference scores on both Stable Diffusion v1.5 and SD3-Medium, notably outperforming the official 28-step baseline with only 20 steps

HONORS & AWARDS

- **Edmund J. James Scholar**, UIUC 2026
- **Distinguished Undergraduate Researcher**, UIUC 2026
- **Dean's List** (all semesters), UIUC 2024 – 2026
- **First Place in 2025 Illinois Datathon** (1/180) 2025
- **Second Place in 2024 Sandia National Lab Data Challenge** (2/51) 2024
- **National First Prize in University Mathematical Modeling Contest** 2023

SKILLS

- **Programming & Frameworks:** Python, PyTorch, $\LaTeX$, Linux, Git, HPC/Slurm
- **Languages:** Mandarin (Native), English (Fluent, TOEFL 105/120)