Mufan Qiu

mufan@cs.unc.edu • Homepage • Github

RESEARCH INTEREST

AI for Science and Efficient ML

- Improving scaling efficiency (via MoE, merging, *etc.*) in LLM and Multimodal LLM.
- Building reliable multi LLM agent system for general and function calling applications
- AI for Science (via RNA foundation model, multimodal modeling, *etc.*).

EDUCATION

University of North Carolina at Chapel Hill, Chapel Hill, NC, USA — Aug 2024 – Present

- Ph.D. in Computer Science
  - Advisor: Prof. Tianlong Chen

University of Science and Technology of China, Anhui, China — Sep 2020 – Jun 2024

- B.S. in Computational Mathematics (GPA: 3.77/4.3)

PUBLICATIONS

(*: Equal contribution)

CONFERENCES

KeyTAR: Practical Keystroke Timing Attacks and Input Reconstruction
Mufan Qiu*, Lihsuan Chuang*, Dohhyun Kim, Huaizhi Qu, Tianlong Chen, Andrew Kwong
IEEE S&P 2026

Bag of Tricks for Sparse Mixture-of-Experts: A Benchmark Across Reasoning, Efficiency, and Safety
Mufan Qiu*, Zheyu Shen*, Pingzhi Li, Ang Li, Tianlong Chen
Findings of EMNLP 2025

`GRNFormer`: A Biologically-Guided Framework for Integrating Gene Regulatory Networks into RNA Foundation Models
Mufan Qiu, Xinyu Hu, Fengwei Zhan, Sukwon Yun, Jie Peng, Ruichen Zhang, Bhavya Kailkhura, Jiekun Yang, Tianlong Chen
➤ In this work, we introduce `GRNFormer`, integrating multi-scale gene regulatory networks into RNA foundation models through biologically-guided hierarchical GRN construction and adaptive structure fusion, achieving 3.6% PCC improvement in drug response prediction and 9.6% AUC gain in single-cell classification while revealing interpretable regulatory patterns.
Findings of ACL 2025

Advancing MoE Efficiency: A Collaboration-Constrained Routing (`C2R`) Strategy for Better Expert Parallelism Design
Mohan Zhang*, Pingzhi Li*, Jie Peng, **Mufan Qiu**, Tianlong Chen
NAACL 2025, Senior AC Award

JOURNALS

Single-cell epigenomics and transcriptomics reveals altered circuitry in Alzheimer's disease across multiple brain regions
Zunpeng Liu, Benjamin T. James, Shanshan Zhang, Kyriaki Galani, Riley J. Mangan, Stuart Benjamin Fass, Sukwon Yun, Carles A. Boix, Yena Sung, Xushen Xiong, Na Sun, Lei Hou, Martin Wohlwend, Yosuke Tanigawa, **Mufan Qiu**, Xikun Han, Lei Xiong, Efthalia Preka, Lei Huang, Li-Lun Ho, Julio Mantero, Hansruedi Mathys, Tianlong Chen, David A. Bennett, Li-Huei Tsai, Manolis Kellis
Cell

IN SUBMISSION/ PREPARING

(*: Equal contribution)

CONFERENCES

SPATIA: Multimodal Model for Prediction and Generation of Spatial Cell Phenotypes
Zhenglun Kong*, **Mufan Qiu***, John Boesen, Xiang Lin, Sukwon Yun, Tianlong Chen, Manolis Kellis, Marinka Zitnik
➤ In this work, we introduce SPATIA, a multi-scale generative and predictive model for spatial transcriptomics. SPATIA learns cell-level embeddings by fusing image-derived morphological tokens and transcriptomic vector tokens using cross-attention and then aggregates them at niche and tissue levels using transformer modules to capture spatial dependencies.
Submitted to ICLR 2026

`scPaLM`: Pre-training of Single-cell Language Models through Genetic Pathway Learning
Mufan Qiu*, Xuxi Chen*, Sukwon Yun*, Zhangyang Wang, Marinka Zitnik, Manolis Kellis, Tianlong Chen
➤ In this work, we present `scPaLM`, a single-cell foundation model that integrates genetic pathway learning through vector quantization, progressive cell token aggregation and permutation-invariant gene embedding, achieving 10.1% improvement in cell annotation accuracy and 5.15% higher drug response correlation while enabling interpretable pathway discovery through discrete codebook learning.
Submitted to ICLR 2026

Symbiotic Cooperation for Web Agents: Harnessing Complementary Strengths of Large and Small LLMs
Ruichen Zhang*, **Mufan Qiu***, Zhen Tan*, Mohan Zhang, Xiaopeng Lu, Jie Peng, Kaidi Xu, Leandro Z. Agudelo, Peter Zhenghao Qian, Tianlong Chen
Submitted to ICLR 2026

RESEARCH EXPERIENCES

Amazon, Seattle, WA — May 2025 – Apr 2026 (Expected)
- Applied Scientist Intern, mentored by Dr. Mostafa Karimi
  - Leveraged Megatron for large-scale pre-training of protein language models.
  - Responsible for the migration from standard huggingface framework to Megatron pre-training framework.
  - Explored Mixture-of-Experts (MoE) and asymmetric architectures for protein modeling.
  - Conducted comprehensive validation of scaling laws for protein language models.

The University of North Carolina at Chapel Hill, Chapel Hill, NC — Feb 2024 – Jun 2024
- Remote Research Intern, mentored by Prof. Tianlong Chen

Microsoft Research Asia, Beijing, China — Oct 2023 – Jul 2024
- Algorithm Engineer Intern in MSRA AI for Science, mentored by Dr. Shuxin Zheng

AWARDS & SCHOLARSHIPS

Awards
- Outstanding Graduate of University of Science and Technology — 2024

Scholarships
- Gold Scholarship for Outstanding Students of USTC — 2023

PROFESSIONAL SERVICES

Reviews
- Conference on Neural Information Processing Systems (NeurIPS) — 2024
- Conference on Artificial Intelligence and Statistics (AISTATS) — 2024
- Workshop on Large Language Models and Generative AI for Health at AAAI — 2025

Tutorials
- International Conference on Machine Learning (ICML) — 2024