

QIANYU ZHENG

Atlanta, GA · zhengqianyu0227@gmail.com · +1 (404) 960-4914
qzheng75.github.io · linkedin.com/in/zheng-qianyu-365b60250

EDUCATION

Georgia Institute of Technology — *M.S. Computer Science, Atlanta, GA* · GPA 4.00/4.00 Aug 2025 – May 2027 (expected)
Georgia Institute of Technology — *B.S. Computer Science, Atlanta, GA* · GPA 4.00/4.00, Faculty Honors Aug 2022 – May 2025

TECHNICAL SKILLS

Languages: Python, Java, SQL, Bash

ML & AI: PyTorch, scikit-learn, LLM agents and multi-agent systems, MCP, graph neural networks, NLP, contrastive learning, hyperparameter optimization

Data & Compute: Apache Spark (PySpark), Tableau, Matplotlib, Seaborn, PowerBI, Pandas, NumPy, NetCDF, MongoDB, high-performance computing (SLURM)

Cloud & Tools: AWS (ECS, Fargate, Lambda, ECR, S3, Route 53), Docker, Flask, Git, Linux

EXPERIENCE

Lila Sciences, Inc. — *Machine Learning Intern, Physical Sciences Simulation Team, Cambridge, MA* May 2026 – Aug 2026

- Shipped an agentic hyperparameter-optimization framework (Python, PyTorch) into internal research infrastructure; tuned models set internal SOTA benchmarks and reached production use.
- Built a multi-agent auto-research system that proposes and evaluates new model architectures, generating proposals concurrently with in-flight training runs to raise experiment throughput over sequential loops.
- Validated it by recovering strong-baseline performance from an underperforming architecture and extending a model to a property outside its original scope.

Georgia Institute of Technology — *Graduate Research Assistant, Fung Group, Atlanta, GA* May 2023 – Present

- Built a differentiable geometry toolbox exposing structural properties as autograd-optimizable losses, served to LLM agents over MCP; 87.5% success on 40 natural-language editing tasks, beating Claude Code and toolbox-free baselines.
- Designed a propose-act-feedback multi-agent system for inverse design, succeeding on 9/10 tasks requiring 3x increases in a target material property.
- Developed a physics-informed pretraining pipeline for ML potentials, validated across 3 datasets and 3 architectures (CGCNN, TorchMD, M3GNet); cut simulation failure modes vs. regularization baselines. First-author JCI paper.
- Building an self-evolving agentic AutoML pipeline to find interpretable predictive models for superconductors with symbolic regression (SISSO).

Fraunhofer Institute for Wind Energy Systems — *Data Engineering Research Intern, Bremen, Germany* Aug 2025 – Dec 2025

- Built PySpark and NetCDF pipelines processing a year of TB-scale multi-instrument sensor data (dual-Doppler radar, lidar, disdrometer, met stations) into a validated benchmark dataset.
- Designed a local-contrast anomaly filter to flag unphysical spikes in radar measurements; confirmed flagged cases showed measurably degraded accuracy against reference instruments.
- Benchmarked two wake models across a 100+ km², two-site domain using the cleaned dataset. First-author JRSE paper.

Leibniz Institute of Plant Biochemistry — *Machine Learning Intern, Halle, Germany* May 2024 – Aug 2024

- Modeled large protein families as sequence-similarity graphs, applying sparsification, community detection, and centrality metrics to rank functionally important sequences.
- Designed a cluster-based data-splitting strategy (ESM embeddings + surrogate-model evolutionary search) that outperformed random splits and SpanSeq on downstream classification.

Georgia Institute of Technology — *Teaching Assistant, CS 1331 Object-Oriented Programming* Jan 2023 – May 2024

- Supported ~500 students per semester in Java and JavaFX across four semesters; served as Q&A forum lead coordinating a 25-TA response team.

PROJECTS

NL2Prot — *Natural-Language Retrieval over Protein Databases · Project Lead (2-person team)* Sep 2024 – Jan 2025

- Trained a CLIP-style dual encoder (ESM2 + BERT) on a 10k-pair corpus to retrieve UniProt entries from free-text queries.
- Reached 94.05% top-50 accuracy on queries vs. 67.33% for an NER/keyword baseline, cutting query latency from over 24 hours to under 40 seconds.
- Deployed a quantized inference service with Flask and Docker on AWS (ECS/Fargate, ECR, Route 53), live at nl2prot.org.

PUBLICATIONS & CERTIFICATIONS

Publications: Zheng, Q. & Fung, V., Journal of Chemical Information and Modeling (2026) · Zheng, Q. et al., Journal of Renewable and Sustainable Energy 18(4) (2026). Both first author.

Awards & Certs: AWS Certified Machine Learning – Specialty · AWS Certified Cloud Practitioner