Adrian Cheung

(571) 888 5457 • acheung@gatech.edu • Ashburn, VA • US Citizen
linkedin.com/in/acheunggt • github.com/adrian-cheung • adriancheu.ng

Education

Georgia Institute of Technology

M.S. in Computer Science, Specialization in Machine Learning • GPA: 4.0/4.0 Aug 2025 – May 2026

Coursework: Deep Learning, Computer Vision, NLP, ML Security, Graduate Algorithms

B.S. in Computer Science, Specializations in AI & Theoretical CS • GPA: 4.0/4.0 Aug 2022 – May 2025

Coursework: ML, Robotics, AI, Algorithms, Data Structures, OOP, Objects & Design (Agile/Scrum), Automata, Computer Architecture, Graph Theory, Adv. Linear Algebra, Statistics, Combinatorics, Discrete Math

Experience

Amazon Arlington, VA

Software Development Intern, FPDS Associate Data Store Team *May 2025 – Present*

- Developed an enhancement for employee data segmentation system to fix bugs impacting 20M+ employees
- Built MCP server with RAG and multi-agent orchestration to streamline support, onboarding, and dev processes

Software Development Intern, AWS Systems Manager Agent Team *May – Aug 2024*

- Created two AWS resource management agents in Go and Rust to investigate migration of a codebase to Rust
- Implemented telemetry, socket-based IPC, and communication with AWS services (CloudWatch, S3, SSM)
- Reduced CPU usage by 20%, memory by 34%, binary size by 30%, high-traffic IPC latency by 83% with Rust

Graph Computation and Machine Learning Lab Atlanta, GA

Undergraduate Research Assistant *Feb 2025 – Present*

- Led project on machine unlearning to defend against data reconstruction and extraction under Dr. Eli Chien
- Proposed evaluation framework for unlearning from the novel perspective of model inversion attacks

Physically-Based Generative Neural Graphics Group Atlanta, GA

Undergraduate Researcher *Jan 2024 – May 2025*

- Investigated video diffusion with transformers and state space models for physics simulations under Dr. Bo Zhu
- Conducted experiments to generate trajectories for heat transfer, reaction-diffusion, and n-body problem

JHU Applied Physics Laboratory Laurel, MD

Software Engineering Intern, Cyber/ML *Jun – Nov 2023*

- Rewrote pipeline to reduce data generation time by 94% and model code by 82% across 14 datasets, 500+ fields
- Prototyped assistant to suggest code from JSONs and natural language requirements using LangChain and LLMs

Publications

Tile-Based ViT Inference with Visual-Cluster Priors for Zero-Shot Multi-Species Plant ID | *CLEF 2025*

- Identified plants from vegetation plots for PlantCLEF 2025, featured on PyTorch webinars and at CVPR 2025
- Performed efficient DINOv2 feature extraction and nearest-neighbor search on 1.4M+ images of 7,800 plant species
- Obtained 2nd highest score with novel tiling, SAM-based segmentation, and prediction aggregation methods

Distilling Spectrograms into Tokens: Fast and Lightweight Bioacoustic Classification | *CLEF 2025*

- Established foundation for fast sequence models in bioacoustics with 120x real-time inference in BirdCLEF+ 2025
- Explored spectrogram clustering tokenization, skip-gram embeddings, and student-teacher pretraining

Transfer Learning with Pseudo Multi-Label Birdcall Classification | *CLEF 2024*

- Developed data pipeline using SOTA model embeddings to identify bird species from recordings
- Earned BirdCLEF 2024 Best Paper and achieved 90% of top score with signficantly less data and training

Skills

Languages: Python, Rust, Go, Java, C#, C++, C, Julia, JavaScript, Bash, HTML/CSS
Frameworks: NumPy, PyTorch, pandas, Matplotlib, Seaborn, Tensorflow, Keras, OpenCV, scikit-learn, SciPy, Luigi, Spark, Petastorm, Flask, SQLAlchemy, Django, Tokio, React Native, JavaFX, JUnit
Tools: Git, Linux, Docker, AWS, Google Cloud, Jupyter, PostgreSQL, Elasticsearch, Bitbucket