

JOSHUA C. HU

US/Canadian Citizen | (315) 418-9365 | jchu24@buffalo.edu | linkedin.com/in/joshua-cong-hu | https://github.com/JoshuaCongHu

EDUCATION

University at Buffalo

B.S. in Computer Science, Concentration: Intelligence, Modeling & Simulation

Buffalo, NY

Expected May 2027

University of Toronto

Bioinformatics & Data Science

Toronto, ON

Sep 2023 – May 2024

EXPERIENCE

AI Fellow (Incoming)

May 2026 – Apr 2027

Break Through Tech, Cornell University

Remote

- Selected for a year-long AI Fellowship with **Cornell Tech**: intensive ML Foundations training plus a team industry challenge project with companies including **Google**, **Snowflake**, and **KPMG**.

Software Engineering Intern

May 2026 – Present

Linde

Buffalo, NY

- Redirected reporting platform from a custom CRUD stack to Microsoft Fabric/Power BI by evaluating 4 architecture options against stakeholder-prioritized criteria, eliminating **~60–70%** of planned custom development and positioning the platform for AI-agent integration via Fabric-tagged data with LLM-driven reporting as a forward path.
- Designed a strangler-fig migration from an aging on-prem SQL row-store to hybrid KQL/Eventhouse + SQL on Microsoft Fabric for **30,000+** IoT devices on a **3.17B-row** telemetry workload; columnar ingest eliminates burst lock contention during mass RTU reconnection and delivers **5–8× storage compression**.

LLM Researcher Intern

May 2026 – Present

Algoverse

Remote

- Built an **LLM evaluation framework** scoring multi-agent tool-calling on quality-vs-cost across **3 axes** (task difficulty, agent heterogeneity, interaction structure); ran GRPO experiments quantifying the cost-accuracy trade-offs behind agent selection (paper under review).

Software Tooling Developer

Jan 2026 – Present

University at Buffalo, CSE Department

Buffalo, NY

- Refactoring a multi-institution code-tracing platform under guidance of a **Vercel** engineering lead; architecting a React + TypeScript + PostgreSQL stack with a Docker deployment pipeline for cross-university adoption.
- Scaling the platform from **800+ students** and **15,000+ submissions** to multi-institution use, with **1 external university** in late-stage adoption for Fall 2026 rollout and system outputs positioned for academic publication.

AWS AI & ML Scholars — Agent Developer Track

Aug 2025 – Nov 2025

Amazon Web Services

Remote

- Selected for the fully-funded program as a **top ~3%** performer (**top 1,500 of 50,000+** participants) based on technical assessment performance.
- Built a health-claims triage agent on **Amazon Bedrock** with an **S3 + Lambda** ingestion pipeline, FAISS vector retrieval, and **guardrails** that raised the compliance-check pass rate **~25%** over an unguarded baseline.

Teaching Assistant

Aug 2025 – Present

University at Buffalo

Buffalo, NY

- Supported **453+ students** across labs, office hours, and grading; ran structured debugging workshops that lifted mentee exam scores **~20% above class median**, and built tailored materials for non-native speakers and students with disabilities that cut assignment completion time **30%**.

LLM Researcher Intern

May 2025 – Sep 2025

Algoverse | Columbia, Meta FAIR, UofT, St. Jude Children's Research Hospital

Remote

- Co-authored a **NeurIPS 2025 Workshop** paper (LoRA-Guided PPO, second author) on **cost-aware agent orchestration**; built a controller that decides which sub-agents to invoke under a budget, cutting **inference cost-per-success 12%** vs. a PPO baseline (**24.30 → 21.40**) while holding **100% task success** on ToolBench.
- Engineered a DAG-based multi-agent orchestrator with on-demand subagents and vector-DB memory; parallel RAG workflows handled tasks **3× larger** than a single context window while cutting **token cost 28%**.

PROJECTS

StructStudio | React, JavaScript, Canvas API, Google Analytics, Playwright

Feb 2026 – Present

- Co-founded an interactive code-visualization platform as an inaugural UB Forge fellow (**1 of 14** from **250+** applicants); renders real-time simulations of Trees, Heaps, and Graphs from raw program input, shipped to production and adopted across theory-level CS courses at the University at Buffalo, including CSE 250 (Data Structures & Algorithms).
- Engineered a React rendering pipeline achieving **sub-100ms render times** while sustaining **300+ concurrent users**; instrumented with Google Analytics for usage telemetry and Playwright for end-to-end test coverage.

LEADERSHIP

Data Science & Machine Learning Club — Founder & President

Jun 2025 – Present

University at Buffalo

Buffalo, NY

- Founded and scaled the club to **120+ members**; secured partnerships with the UB AI/ML department to source speakers from **Vercel**, **Amazon**, and **Google**, and built real-world project briefs for **6+ technical workshops**.
- Managed a **6-person executive team** and **\$1.5K+ annual budget**, structuring operations across event logistics, outreach, and content planning to sustain weekly programming.

TECHNICAL SKILLS

Languages: Python, Java, JavaScript, TypeScript, SQL, KQL, C, HTML, CSS

AI / ML: PyTorch, TensorFlow, scikit-learn, LangChain, FAISS, pandas, NumPy, Embeddings, RAG, LoRA, PPO

Web & Frameworks: React.js, Next.js, Node.js, Express.js, REST APIs, WebSockets

Tools & Platforms: Git, Docker, AWS (S3, EC2, Lambda, SageMaker, Bedrock), PostgreSQL, Vercel, Microsoft Fabric

Concepts: Multi-Agent Systems, Agent Orchestration, Evaluation Frameworks, AI Guardrails, System Design, Vector Databases