

Austin Senna Wijaya

646-295-5341 | asw2215@columbia.edu | linkedin.com/in/austin-senna | github.com/Austin-Senna | austin-senna.github.io

EDUCATION

Columbia University

B.S. in Computer Science

New York, NY

Aug 2024 – May 2028

- GPA: 4.22/4.0; Full-Tuition Scholarship; Dean's List
- Coursework: Deep Learning, Natural Language Processing, Databases, Operating Systems, Distributed Systems

EXPERIENCE

Software Engineer Intern

Jun 2026 – Aug 2026

Garner Health

New York, NY

- Deployed a multi-agent on-call triage system on **Kubernetes** that autonomously raises alerts from telemetry, investigates root cause, and ships fix pull requests, built with **Claude Agent SDK** and custom MCP servers for PostgreSQL, GitLab, Grafana; resolved **15+ production incidents**, saving on-call engineers **10+ hrs/week**.
- Built an observability stack for a B2B web portal with **OpenTelemetry** and **Datadog**, adding distributed tracing and metrics across frontend and backend microservices that surfaced **40+ user-impacting bugs** and **tripled** the team's bug resolution speed.

AI Research Assistant

Dec 2025 – Present

Columbia Data, Agents, and Processes Lab

New York, NY

- Engineered a benchmark evaluation harness in **Strands SDK** that runs question-answering agents as parallel worker processes, each with dataset-retrieval and sandboxed **Python/DuckDB** code-execution tools, making evaluation **4x faster**.
- Architected telemetry to track tool-call and reasoning traces to pinpoint agent failure bottlenecks; resulting harness fixes raised GPT-5-mini benchmark accuracy from **5.16%** to **56.3%**.

Machine Learning Research Assistant

Dec 2025 – May 2026

Columbia Zuckerman Mind Brain Behavior Institute

New York, NY

- Built and parallelized a GPU-optimized **PyTorch** pipeline in HPC clusters for **Qwen2-Audio** inference, hidden-state extraction, and supervised classification training under cross-validation across 12 speech attributes and **120K+ audios**.
- Trained bottleneck FNNs to compress **4096-D** hidden-state embeddings into **10-100D** feature-specific representations, isolating individual linguistic features from **Qwen2-Audio** for downstream analysis.

PUBLICATIONS

- SANA: What Matters for QA Agents over Massive Data Lakes?* **VLDB DASHSys 2026**. First author.
- MINA: Linear Probes Reveal Coding-Sequence Family Signal in Frozen DNA Encoders*. **MLCB 2026**. First author.
- LakeQA: A Benchmark for Complex Exploratory QA over a Million-Scale Data Lake*. **ICML 2026**. Fourth author.
- Targeted Embeddings: Neural Speech Encoding with Large Audio-Language Models*, **NeurIPS 2026**: Third author (*Under-review*).

PROJECTS

ResearcherX: Local-First AI Research IDE | *Next.js, FastAPI, LanceDB, NetworkX, Electron*

[Website](#)

- Led development of a local-first AI research editor using **GraphRAG** to flag unsupported claims, contradictions, and citation drift in AI-assisted drafts; **2nd Place** in **Columbia's flagship startup competition**.
- Engineered a real-time contradiction-detection engine that re-analyzes only edited paragraphs via content-hash caching and batches up to 100 claims per Pydantic-validated LLM call, cutting LLM costs by **90%**.

TECHNICAL SKILLS

- Languages:** Python, Java, Golang, C, JavaScript / TypeScript, SQL
- AI & Data:** PyTorch, Hugging Face, Strands Agents, Claude Agent SDK, Langchain, Pandas, Playwright, LanceDB
- Infra & Cloud:** Kubernetes, Docker, Linux, Git, Terraform, AWS (EC2, S3, Bedrock), GCP, Grafana, OpenTelemetry
- Framework & Tools:** Next.js, React, FastAPI, Litestar, Electron, Claude Code

HONORS

- Jane Street Mystery Planet Trading Competition:** 1st Place.
- Databricks Grounded Reasoning Cup:** Top 10 representing in Columbia mentored by Google team.
- Hack@Brown:** Finalist; Strongest Product Thinking award; worked with Google Ventures and Partiful.
- Tencent Kepler Plan Hackathon:** 2nd Place; Best Code; Best Product Idea.
- 55th International Chemistry Olympiad:** Final stage, 10th rank for Indonesian national team selection.