

# Alexander Liu

aliksliu@gmail.com | alexanderzliu.com | linkedin.com/in/alexanderzliu | github.com/alexanderzliu

## Skills

---

**Languages:** Python, SQL, TypeScript

**AI Developer Tooling:** Claude Code, OpenAI Codex, MCP, skills, subagents, LLM observability & evals

**Orchestration & Services:** Temporal Python SDK, Airflow, FastAPI, GitLab CI/CDwhat

**Data, ML & Infrastructure:** Databricks, Snowflake, dbt, MLflow, Git, AWS Bedrock/EC2

## Experience

---

### AI Solutions Engineer

Oct 2025 – Present

Penguin Random House

San Francisco, CA (Remote)

- Own and operate PRH's internal Claude Code plugin marketplace hosting MCP servers, skills, and subagents along with command-line scripts for Claude Code set up and plugin onboarding across a 20+ person data science org.
- Drove agentic coding tool adoption through office hours, workshops, team residencies, tutorial repos, docs, and plugin support; team survey showed 19/21 used daily, 20/21 used plugins and 20/21 felt help was easy to access.
- Built feedback loops for AI developer productivity, combining cost monitoring dashboards, SPACE-inspired surveys, and qualitative research to identify friction and guide plugin marketplace improvements.
- Established engineering practices for production-grade internal agents, including observability, evals, deterministic validation, and cost controls; used them to ship anti-piracy and metadata agents covering 40k+ YouTube videos, \$8M+ revenue at risk, and 50k+ authors, with 8x lower token spend and >99% eval accuracy.

### Machine Learning Scientist

Apr 2024 – Oct 2025

Penguin Random House

San Francisco, CA (Remote)

- Built a reusable orchestration library for data science projects, with templates for Airflow DAGs and Databricks jobs plus GitLab CI/CD deployment scripts, standardizing workflow deployment across the org.
- Led production deployment of a new-title pricing pipeline supporting 2,400+ books/year, composing Databricks training/inference workflows with automated data-drift detection and alerting.
- Reduced demand forecast error up to 39% on 40k+ titles weekly by fine-tuning pretrained transformer time-series models (Chronos), validated via historical backtests.

### Project Management Analyst

Oct 2021 – Apr 2024

Penguin Random House

San Francisco, CA (Remote)

- Stood up PRH's first cross-functional GenAI program with 40+ members, prioritizing use cases, co-authoring the enterprise AI risk framework, and moving early production cases from POC to rollout.
- Led technical evaluations for 10+ AI vendors across coding, copywriting, and document understanding, designing scoring frameworks and aligning Legal, InfoSec, and Editorial on procurement decisions.

## Projects

---

### Porygon – Durable Agent Harness & Eval Framework | [github.com/alexanderzliu/porygon](https://github.com/alexanderzliu/porygon)

- Built a Temporal-orchestrated harness for a long-horizon agent playing Pokémon Red. Trials are durable workflows; agent's loop runs as idempotent activities with parent/child fan-out for parallel experimentation.
- Created the eval & observability framework on PyBoy: programmatic success conditions over gamestate, per-step cost and behavior tracking, and inspectable per-trial artifacts to enable replay and post-hoc analysis.
- Implemented a modular agent harness with swappable prompts, game state representation and memory-strategy modules to enable experimentation with different agent architectures.

### Crimson | [github.com/alexanderzliu/aws-dd-hackathon](https://github.com/alexanderzliu/aws-dd-hackathon) | Winner at AWS LLM Observability Hack-a-Thon

- Built a multi-agent red-teaming platform (AWS Strands Agents SDK) that maps a target agent's architecture to a Neo4j graph and runs adversarial multi-turn conversations to surface prompt injection vulnerabilities.
- Instrumented the pipeline end-to-end with Datadog LLM Observability and shipped a real-time FastAPI + neovis.js dashboard with per-attack traces so that attacks evolve with the target agent's architecture.

## Education

---

U.C. Berkeley – B.A. in Computer Science

Aug 2021