

Alex Toohey

415-876-1905 | San Francisco, CA | alex.toohey@gmail.com | [LinkedIn](#) | [GitHub](#) | alextoohey.dev

EDUCATION

University of California, Berkeley

May 2028

Bachelor's Degree in Data Science (Economics Domain Emphasis), Minor in Computer Science

GPA: 3.93/4.0

Relevant Coursework: Object-Oriented Programming, Data Structures and Algorithms, Data Engineering, Optimization, Machine Learning, Linear Algebra, Probability, Discrete Math, NLP, Computer Vision, Time Series, R

Activities and Leadership: Data Science Society (Project Manager), Code for Good (Recruitment Chair), Sports Analytics Group (Data Journalism Editor), Data & CS Scholars, Data 8 Course Staff, Data Discovery, Berkeley SkyDeck

EXPERIENCE

Visa

May 2026 – Present

Data Science Intern

Foster City, CA

- Engineered a distributed Spark ETL pipeline on Kubernetes to preprocess **76+ billion** raw payment transactions into Hive training datasets with natural-language feature representations tailored for scalable LLM fine-tuning
- Developed an LLM-based merchant recommendation model, extending Llama with **10,000+** new merchant tokens, domain-specific embedding initialization, and LoRA fine-tuning, achieving **84%** Recall@10 and **88%** NDCG@10
- Executed **210+** ablation studies optimizing data representations, prompt templates, and training configurations, increasing previously unseen merchant recommendation performance by **6×** over strongest heuristic baseline

Databricks

March 2026 – June 2026

Software Engineer (Contract)

Berkeley, CA

- Built scalable testing infrastructure for Databricks' Partner Demo Catalog with pytest, Playwright, and Databricks SDK, automating validation across **169** enterprise AI/ML partner solutions and reducing verification time by **94%**
- Designed a modular two-layer validation architecture combining browser automation with extensible backend validation, ensuring end-to-end coverage of user flows and **12+** Databricks resource types through **350+** tests
- Developed an automated reporting pipeline streaming validation results to a Delta table after every test run, enabling historical analysis, CI/CD, and early detection of **36+** production issues before partner release

Bay Area Rapid Transit (BART)

January 2026 – May 2026

Software Engineer Intern

Berkeley, CA

- Built a secure, fully local code intelligence platform integrating static analysis, code graphs, and grounded LLMs to generate architecture documentation and diagrams for **20,000+** LOC production codebases in **under 2 minutes**
- Engineered a modular parsing framework leveraging Tree-sitter and ASTs to support **8+** programming languages
- Developed an LLM pipeline with static grounding and multi-stage validation, achieving **97.5%** summary accuracy

SubscriptionFlow (YC W22)

June 2025 – August 2025

AI/ML Engineer Intern

London, UK

- Built a RAG-powered Text-to-SQL model with VannaAI, GPT-4o, and ChromaDB, cutting retrieval time by **92%**
- Embedded **500+** SQL Q&A pairs, **140+** schemas, and **55+** docs in a vector DB, reducing hallucinations by **75%**
- Designed a 4-class customer churn prediction pipeline with AWS SageMaker, enabling targeted customer retention

PROJECTS

AskOski: ML & Data Infrastructure | *Python, Docker, PyTorch, Airflow*

September 2025 – May 2026

- Engineered Airflow ETL pipelines to ingest, transform, and store anonymized academic records in PostgreSQL, retraining five PyTorch models deployed across AskOski, a Berkeley advising platform for **45,000+ students**
- Refactored a BERT course recommender, cutting real-time personalized inference latency to under **250 ms**
- Created multi-stage Docker builds with cache-free installs, cutting pipeline image size **80%** from 25 GB to 5 GB

SKILLS AND INTERESTS

Programming Languages: Python, Java, R, SQL (MySQL, Postgres), HTML/CSS, JavaScript/TypeScript (Node.js)

Libraries & Frameworks: PyTorch, Hugging Face, Flask, FastAPI, React, Spring Boot, pytest, Playwright, Next.js

Cloud & Big Data: GCP, AWS, Snowflake, Apache Spark, Hadoop, Hive, Airflow, Trino, Firebase, Redis

Software & Tools: Git, Docker, Kubernetes, GitHub Actions, Figma, Jira, Claude Code, Codex

Interests: Tennis, Skiing, Hiking, Piano, Mandarin, Baking, Reading