

Oliver Carmont

ocarmont@berkeley.edu | linkedin.com/in/oliver-carmont | github.com/olivercarmont

EDUCATION

University of California, Berkeley

Bachelor of Arts in Computer Science

Berkeley, CA

Aug 2021 – May 2025

- Coursework: Data Structures, Efficient Algorithms, Principles of Data Science, Machine Structures, Discrete Math, Introduction to Artificial Intelligence, Computer Security, Operating Systems
- Activities: President of Cal Hacks, one of the largest collegiate hackathons nationally.

EXPERIENCE

Apple

ML Software Engineer Intern

May – Aug 2024

SF, CA

- Undisclosed AI/ML project at Apple's Security & Privacy team for automating security reviews.
- Involving fine-tuning, evaluations, RAG (Retrieval Augmented Generation), and agents.
- Worked with network security protocols and compliance frameworks (ISO, SOX, and PCI)
- Gave a presentation to the VP of Engineering to secure future funding for this project.
- Secured 3rd place at Apple's internal hackathon with 360+ participants, for optimizing LLM model-size routing.
- Contributed to internal AI/ML libraries for agentic functions and model wrappers.

Amazon

Software Engineer Intern

May – Aug 2023

Sunnyvale, CA

- Architected a data pipeline integrating AWS Lambda with Docker, Cloudformation, Quicksight, and Redshift, powering performance analytics for Amazon Alexa. Worked with CI/CD pipelines, code reviews and pushed to prod.
- Delved into audio technology and firmware concepts, including DSP & AFE algorithms for echo cancellation.
- Collaborated cross-functionally with audio scientists, and AWS experts, to save 150 employees, 1 hour every week.

UC Berkeley College of Engineering

Student Course Lecturer

Sep 2022 – Present

Berkeley, CA

- Lecturer for CS 198: Introduction to Web Development, a UC Berkeley course run by students.
- In a team of 4, I teach lectures and discussions, and grade homeworks for 85 students. Learn more at cubstart.com

RESEARCH

Berkeley Sky Computing Lab

Undergraduate Researcher

Sep 2024 – Present

Berkeley, CA

- Working with PhD student, Asim Biswal, on Natural Language-to-SQL frameworks and benchmarks for LLMs to query structured data. Working in a lab alongside the founders of Databricks and Anyscale.

Berkeley Artificial Intelligence Research (BAIR)

Undergraduate Researcher

Jan – Jun 2024

Berkeley, CA

- Undergraduate research, supervised by PhD student, Akshat Gupta, at BAIR's Speech Group, to fine-tune a model with LoRA and create a RAG workflow, to mimic the conversational style of a user.

PROJECTS

(Side Project) Financial Data Visualizer AI | C++, Ruby on Rails, Sagemaker, EKS

Aug 2024 - Present

- Building a platform for generating financial graphs from natural language (using text-to-SQL-to-visual). Fine-tuned and evaluated on text-to-SQL benchmarks like WikiSQL. AWS Sagemaker for fine-tuning and inference. Ruby on Rails web app with Postgres. Hosted on ECS and EKS.

(Class Project) Secure File-Sharing Platform | Go

Apr 2024

- Built in Go from scratch. Wrote a 2-page architectural design document. Integrated: distributed ledgers, cryptographic functions including SHA-512 hashing, symmetric key encryption, public key encryption, and digital signatures.

Carbonly | MongoDB, Express, React, NodeJS, Chrome Dev

Feb – Sep 2020

- Developed a full-stack web app that tracked consumer carbon footprints across 30+ retail sites, powered by a carbon estimation ML model. Given a funding grant of 10,000 euros by the largest renewable energy provider in Finland.

TECHNICAL SKILLS

Languages: JavaScript, Python, Java, C, C++, Rust, Go, Solidity

Frameworks: React, TypeScript, Flask, MongoDB, PostgreSQL, GraphQL

Developer Tools: Docker, Kubernetes, AWS, gRPC, Linux, Hugging Face, PyTorch, Tensorflow

3x AWS Certified Associate-level: Solutions Architect, Developer. Foundational: Cloud Practitioner