

Zach Phelps

Palo Alto, CA | zach@zachphelps.com | zachphelps.com | linkedin.com/in/zachphelps | github.com/zphelps

Education

Stanford University, BS in Computer Science + Artificial Intelligence

Sept 2022 – June 2026

Relevant Coursework

- Data Structures
- Algorithm Analysis
- Machine Learning
- Reinforcement Learning
- Natural Language Processing
- Operating Systems
- Bioinformatics Algorithms
- Data Science
- OOP Systems Design

Experience

Software Engineering Intern (AI Team), SEP – Indianapolis, IN

June 2024 – Aug 2024

- Developed data ingestion and training pipeline for time series ML classification model which identifies anomalies in hydraulic fracking gun perforating systems. Achieved >90% model accuracy.
- Led team to build frontend (React) for an internal LLM sandbox using open-source LLMs hosted on-premises, enabling secure conversations with company and client data.
- Engaged directly with clients during weekly check-ins and technical demos.

Software Developer, Central – Los Angeles/Remote

May 2023 – Sept 2023

- Led development of startup's mobile (Flutter) and web (React) applications which offered field management services tailored to the mobile waste compaction industry.
- Trained image classifier to measure compaction level of industrial waste bins.

Data Science Intern, Honely: Real Estate Predictive Analytics – Remote

June 2022 – Sept 2022

- Developed web-scraping tool to generate sales leads based on criteria such as target geographic region and business type. 1000+ leads generated.
- Collaborated closely with development team, offering targeted feedback to improve UI/UX.

Skills

Languages: Python, Java, C++, HTML/CSS, JavaScript, Dart, SQL

Frameworks & Libraries: React, React Native, Flutter, PyTorch, Next.js, Express.js, Pandas, Numpy, PyTorch

Databases: NoSQL, MySQL, PostgreSQL

Other: Google Cloud Platform, Docker, Agile, Git, Figma

Projects

Formula 1 Podium Prediction | Python, scikit-learn, PyTorch, matplotlib

- Developed ML models (LogReg, RF, XGBoost, NN) to predict F1 podium finishers using historical race data.
- Achieved 54.2% winner prediction accuracy and 0.94 ROC AUC with XGBoost.

PII Redaction with Lightweight Fine-tuned LLMs | Python, PyTorch, Unsloth, Faker

- Developed a privacy-preserving redaction pipeline using synthetic data and fine-tuned BERT and Llama-3.2-1B for detecting and redacting 10 types of PII.
- Achieved 99.3% accuracy with Llama and 97.2% with BERT.

EngScribe: PDF Ingestion and Search Web Tool | Next.js, GCP, TypeScript, PostgreSQL, Stripe

- Developed a web tool to view, merge, tag, and search long (500+ page) PDF documents with both handwritten and typed text, solving a major pain point for VEX Robotics teams.

Get Up Park: Student Engagement Mobile App | Flutter, NoSQL, GCP, Figma

- Developed Flutter app to manage school news, sports scores, club communication, and lunch menus for Park Tudor High School.
- ~250 DAU and >90% student-body adoption.