

Roy Ho

San Francisco, CA | royho.career@gmail.com | [linkedin.com/in/royho1/](https://www.linkedin.com/in/royho1/) | github.com/royho1/ | royho-career.com

EDUCATION

UNIVERSITY OF CALIFORNIA, DAVIS

Bachelor of Science

September 2022 – March 2026

Major in Statistical Data Science and Minor in Computer Science

Relevant Coursework: Statistical Learning, Regression, Probability Theory, Advanced Computing, Algorithm Design & Analysis, Data Structures

Certifications: NVIDIA – Fundamentals of Deep Learning

October 2025

EXPERIENCE

Journal of Artificial Intelligence and Knowledge Engineering (JAIKE)

Artificial Intelligence Researcher

January 2025 – Present

- Conducted research on LLM-based automation, focusing on API-driven system design, agent orchestration frameworks, and productivity applications across research, coding, and enterprise workflows.
- Authored a 25-page research paper on LLM-based automation and agent architectures, synthesizing peer-reviewed and industry research on API integration, architectural design patterns, system limitations, and responsible deployment; submitted for journal publication.
- Served as a peer reviewer for JAIKE, evaluating research on retrieval methods in large language models, reasoning performance in extended tasks, and large-scale model architectures for methodological rigor and evaluation quality.

TechSprint Innovators

Head of Data Engineering

March 2024 – September 2025

- Built a multi-factor stock screening model using fundamental, technical, and NLP-based sentiment features (FinBERT).
- Developed and evaluated a supervised classification model in scikit-learn to predict price appreciation, performing feature selection, model tuning, and performance validation on historical market data.
- Engineered and automated a daily ETL (Extract, Transform, Load) data pipeline (Python, yfinance, Alpaca API) running on a Raspberry Pi to filter equities, generate structured CSV outputs, and deliver real-time investment signals via Discord webhook.

AI Student Collective (AISC)

General Member

September 2024 – April 2025

- Built machine learning models for stroke risk prediction and real-time drowsy driver detection within structured project cycles.
- Performed data preprocessing, feature engineering, model development, evaluation, and project presentations.
- Worked within quarter-long sprint cycles with defined milestones, code reviews, and final project demos, following structured machine learning development workflows from ideation to deployment.

SIGNIFICANT PROJECTS

San Francisco Restaurant Safety Map

April 2026 – May 2026

- Built a full-stack web app mapping 26,000+ health inspections across 6,200+ San Francisco restaurants using the DataSF public inspection feed.
- Engineered a Python ETL pipeline (pandas, geopy) normalizing inspection records into a 3-table SQLite schema with cached geocoding.
- Built a Flask REST API with 6 endpoints and a React + Mapbox frontend featuring debounced search, ZIP-code filtering, geolocation-based "Near Me" centering, and an Insights panel surfacing per-ZIP best/worst rankings via SQL window functions.

Job Market Analytics Dashboard

September 2025 – December 2025

- Owned the data acquisition layer for a job market analytics dashboard, building scrapers with JobSpy and Selenium to collect 9,000+ statistics and data job postings across CA, NY, and TX.
- Built the resume-to-job matching engine using TF-IDF (Term Frequency-Inverse Document Frequency), cleaning and standardizing scraped job descriptions with NLP (spaCy) to extract technical skills, and improving match accuracy from 30% to 60%.
- Delivered the scraped dataset and matching output into the final dashboard, presenting findings to faculty and peers.

Drowsy Driver Detection System

January 2025 – March 2025

- Built a real-time drowsy driver detection system processing live webcam frames to classify open vs. closed eye states, awarded Best Execution for the AISC Winter 2025 cycle.
- Implemented detection with dlib's 68-point facial landmark model, computing Eye Aspect Ratio via scipy Euclidean distance and flagging drowsiness when EAR held below a 0.25 threshold across 20 consecutive frames, plus mouth landmark tracking to catch yawning.
- Trained and evaluated a scikit-learn classifier on the Kaggle MRL eye-state dataset (4,000+ images) to supplement the threshold logic, integrating it into a live OpenCV feed with a pygame audio alert.

SKILLS

Programming: Python, R, SQL, JavaScript, HTML, CSS

Data Analysis & Machine Learning: Excel (advanced formulas, pivot tables, VLOOKUP, data visualization), Regression (Lasso, Ridge, Logistic), Random Forest, Gradient Boosting, CNNs, Feature Engineering, Cross-Validation & Hyperparameter Tuning, Model Evaluation (ROC-AUC, Precision/Recall), TF-IDF, NLP (spaCy)

Libraries & Frameworks: scikit-learn, pandas, NumPy, PyTorch, Flask, SQLAlchemy, SQLite, React, Vite, Mapbox GL JS, BeautifulSoup, Selenium, JobSpy, PySpark, Matplotlib, geopy, scipy, dlib, OpenCV

Tools & Platforms: Tableau, Power BI, Git, GitHub, Jupyter Notebook, VS Code, Cursor, Raspberry Pi