

Cyrus Navasca

415-218-3792 | cyrus.navasca54@gmail.com | linkedin.com/in/cyrusnavasca | cyrusnavasca.com

EDUCATION

University of California, Santa Barbara

Santa Barbara, CA

Bachelor of Science in Statistics & Data Science, Minor in Mathematics

Dec. 2025

- Relevant coursework: Statistical Machine Learning, Advanced Linear Algebra, Time Series Analysis, Probability & Statistics, Applied Stochastic Processes, Bayesian Data Analysis

EXPERIENCE

Data Scientist

Jan. 2026 – Present

Elas

Pleasanton, CA

- Drove \$150k in revenue by architecting a lead generation workflow leveraging XGBoost for prospect scoring, semantic similarity for ranking ICP match and Tavily to synthesize company intelligence for targeted outreach.
- Engineered a time-series anomaly detection pipeline using S-H-ESD and Isolation Forest methods to identify irregular DHCP traffic for 5M+ daily events across 15 locations, reducing false-positive alerts by 35%.
- Developed an agentic time-series forecasting pipeline built on LangGraph to automate model selection, tuning and data retrieval across 15 unique domains, replacing 5+ hours of manual analysis per-series with instant insights.
- Shipped a production-grade file ingestion platform processing 50k+ files by designing scalable REST APIs supporting document upload, text extraction, and durable data persistence between AWS S3 and PostgreSQL.

Data Science Intern

July 2025 – Dec. 2025

Elas

Pleasanton, CA

- Deployed a sales forecasting pipeline integrating Prophet and XGBoost across 10k+ historical transactions, reducing MAPE by 22% through temporal feature engineering and hyperparameter optimization.
- Built an agent evaluation framework to benchmark response quality, retrieval precision and tool selection accuracy over 1k+ test scenarios to identify failure modes, improving RAG faithfulness by 40%.
- Delivered monthly analytics reports to stakeholders by owning KPI dashboards in Tableau covering 15 metrics.

Data Science Research Assistant

Mar. 2025 – Dec. 2025

META Lab

Santa Barbara, CA

- Applied unsupervised learning techniques including PCA and K-means clustering to identify distinct behavioral profiles among 500+ participants, increasing silhouette score from 0.2 to 0.61.
- Built an NLP clustering pipeline using transformer-based embeddings and HDBSCAN on 2,000+ free-response survey entries, surfacing 3 latent thematic clusters that revealed patterns missed by manual annotation.

Data Analytics Intern

May 2025 – Oct. 2025

Daily Nexus

Santa Barbara, CA

- Published a data-driven web article analyzing trends between Rate My Professors ratings and university grade distributions, featuring interactive visualizations that drove dynamic exploration for 30,000+ readers.
- Won 2nd place statewide for Best Interactive Graphic among 15k+ participants at the California Media Awards.

PROJECTS

NFL Play Intelligence

[GitHub](#)

- Built an XGBoost classifier to predict run vs. pass on 8 seasons of NFL play-by-play data, increasing ROC-AUC by 0.31 over baseline using Optuna hyperparameter optimization, feature engineering and stratified cross validation.

Acuity.ai

[GitHub](#)

- Architected an AI-powered clinical assistant, leveraging hybrid-search RAG and gpt-4o-mini to ground readmission assessments from 20+ guideline documents served behind a Dockerized FastAPI service.
- Developed an ensemble modeling pipeline predicting 30-day ED readmission risk on 100,000+ patient records, applying SMOTE to handle class imbalance and achieving 0.93 ROC-AUC.

TECHNICAL SKILLS

Core: Python (NumPy, Pandas, Matplotlib, Scikit-learn), Go, R, SQL

AI/Machine Learning: Supervised Learning, Unsupervised Learning, Deep Learning (PyTorch, HuggingFace), NLP, Time Series, Model Evaluation, Agentic AI (LangChain, FastMCP, FastAPI, Langfuse), Prompt Engineering

Developer Tools: Git, Docker, AWS (S3, SageMaker), PostgreSQL, Jupyter