Cyrus Navasca

Richmond, CA ❖ 415-218-3792 ❖ cyrus.navasca54@gmail.com ❖ linkedin.com/in/cyrusnavasca ❖ cyrusnavasca.com

EDUCATION

University of California, Santa Barbara (UCSB)   **August 2022-June 2026**
Bachelor of Science (B.S.) in Statistics & Data Science, Minor in Mathematics   Santa Barbara, CA
- *Extracurriculars/Awards*: President of UCSB Data Science Collaborative, Dean's Honor Roll

SKILLS

- **Programming:** Python (Scikit-learn, PyTorch, LangChain, MLFlow, ZenML), R, SQL
- **Tools:** AWS, Docker, Git, PostgreSQL, MongoDB, Tableau, Microsoft Excel

EXPERIENCE

Data Science Intern   **July 2025-Present**
Elas   Pleasanton, CA
- Engineered an AI forecasting agent using LangChain to interpret natural language queries from enterprise users and leverage a custom ensemble of Prophet, XGBoost and SARIMA models to deliver insights across key business metrics.
- Integrated Vanna AI with PostgreSQL and AWS with multi-tenant support, enabling clients to convert text-based prompts into optimized SQL queries, automating data retrieval for data-driven decision making.
- Reduced LLM token costs by 65% and response latency by 90% through development of a data pipeline for an AI data visualization agent, extracting metadata from input datasets and optimizing output generation for end users.

Data Science Research Assistant   **March 2025-Present**
META Lab   Santa Barbara, CA
- Constructed a natural language processing workflow for free-response career data, leveraging Sentence Transformers to generate high-dimensional vector embeddings, boosting downstream predictive model accuracy by 58%.
- Designed a segmentation model of over 2,000 questionnaire respondents on cannabis use by applying PCA and K-means clustering to identify "low-use" and "heavy-use" cohorts to classify users into meaningful subgroups.
- Spearheaded statistical analysis of experimental data to uncover relationships among curiosity, creativity, and intervention types using multivariate regression, mediation models, ANCOVA and A/B testing.

Data Analyst Intern   **October 2024-May 2025**
Daily Nexus   Santa Barbara, CA
- Implemented an end-to-end data workflow leveraging Selenium to scrape 3,000+ websites and transforming raw inputs by performing feature engineering and extracting insights from unstructured text to enable robust analytics.
- Published a data-driven web article analyzing trends between Rate My Professors ratings and UCSB grade distributions, embedding interactive visualizations to allow for dynamic data exploration for over 30,000 readers.

PROJECTS

Acuity.ai   GitHub
- Developed an AI clinical assistant that parses medical documents using text extraction and regex-based NLP to structure patient data, and predicts patient readmission risk using an ensemble model with 96% accuracy.
- Architected a retrieval-augmented generation (RAG) pipeline leveraging TF-IDF and FAISS vector search, using similarity-based retrieval to efficiently surface relevant medical knowledge for chatbot responses.
- Deployed the system as an interactive FastAPI service, managing session-specific chatbot states, integrating parsed patient data, model predictions, and retrieved evidence into real-time clinical responses.

Detecting Diabetic Retinopathy with Computer Vision   GitHub
- Built an end-to-end deep learning pipeline in PyTorch by training a ResNet50 CNN, identifying diabetic retinopathy from retinal images with 96% accuracy, enhancing early diagnostic capabilities.
- Leveraged AWS S3 to store and retrieve over 3,000 patient retina scans, enabling seamless access to training data and supporting scalable, cloud-based model development for medical image classification.