

# Adivya Iyer

818-277-6497 | [aditya4iyer@gmail.com](mailto:aditya4iyer@gmail.com) | [linkedin.com/in/aditya-a-iyer](https://linkedin.com/in/aditya-a-iyer) | [github.com/aditya-iyer1](https://github.com/aditya-iyer1)

## EDUCATION

---

- University of California, San Diego** Sep 2025 – Dec 2026  
*M.S., Data Science* San Diego, CA
- **Relevant Courses:** Efficient AI, Recommender Systems, Sequential Decision Making, Stochastic Optimization
- University of California, Santa Barbara** Sep 2022 – Jun 2025  
*B.S., Data Science & Statistics* Santa Barbara, CA
- **Relevant Courses:** Machine Learning, Stochastic Processes, Nonparametric Modeling, Big Data Analytics

## EXPERIENCE

---

- Analyst, Financial Planning & Analysis** Mar 2025 – Sep 2025  
*Allegiant Travel Co.* Remote/Las Vegas, NV
- Built supervised & unsupervised models (Regression, PCA, GMMs) for crew reliability forecasting
  - Designed end-to-end ML pipeline spanning AWS, SQL, and Python
  - Improved out-of-sample forecast accuracy from 25% to within 1% monthly through model iteration
  - Developed LLM-driven assistants via Copilot to automate reporting workflows and explain policy documents
- Research Assistant** Jun 2024 – Aug 2024  
*UC Irvine, ISI-BUDS (SIBS)* Irvine, CA
- Built mixed-effects regression models with multiple imputation to estimate causal effect sizes
  - Applied social network analysis to correct sampling bias in observational survey data
- Backend Developer Intern** Aug 2023 – Jun 2024  
*FirstHive* Remote/Cupertino, CA
- Built predictive models for user churn and retention to support product decision-making
  - Optimized backend performance to reduce API latency and improve analytics responsiveness
  - Collaborated with product & software engineering teams to productionize model outputs
- Finance & Business Intelligence Intern** Jun 2021 – Jul 2023  
*Ecosense Lighting* Los Angeles, CA
- Built SQL-based dashboards to support sales forecasting and performance monitoring
  - Integrated ERP and CRM data sources to improve data quality for downstream analytics

## PROJECTS

---

- ContextCliff CLI** | *Python, LLM APIs, vLLM, SQLite* Sep 2025 – Ongoing
- Developing a CLI tool to measure effective context limits of LLMs using naturally long-context tasks
  - Defining variance-based failure metrics to identify LLM 'context cliffs' where reasoning stability degrades
  - Engineering an automated pipeline to stratify long-context datasets into token deciles for binned evaluation
- Hallucination-Resistant RAG System** | *LangChain, ChromaDB, HuggingFace, Streamlit* Sep 2025 – Ongoing
- Built conservative RAG system prioritizing abstention over hallucination for finance documents
  - Implemented intent-driven hybrid retrieval (semantic + keyword) to improve recall and auditability
  - Logged retrieval scores, grounding checks, and abstention decisions for failure analysis
- GUI Agent Testing** | *LLMs, Agentic AI, AWS* Sep 2025 – Jan 2026
- Evaluated multimodal agentic LLMs on GUI tasks using the OSWorld benchmark
  - Testing agent configurations across text and image observations with structured action spaces
  - Applying prompt engineering to improve task success, robustness, and generalization
- ChanceCalc Stochastic Calculator** | *OpenAI API, Supabase, Selenium, LangChain* Dec 2024 – Jun 2025
- Led development of Python framework automating Stochastic Information Packet JSON generation
  - Combined deterministic metalog fitting with LLM-assisted semantic interpretation and bound inference
  - Reduced SIP generation time from 2-3 hours to <5 minutes while preserving guardrails

## TECHNICAL SKILLS

---

**Programming:** Python, SQL, R, Java, JavaScript, SAS, HTML/CSS  
**AI/ML:** PyTorch, TensorFlow, Scikit-learn, LangChain, vLLM, HuggingFace  
**Frameworks & Libraries:** PySpark, Databricks, GCP, SQL, AWS