

Aryan Agarwal

9867083601 • aryanagarwal1227@gmail.com • www.linkedin.com/in/aryan1227 • aryan-agarwal-a70cd.web.app

EDUCATION

California State University, Fullerton, California

Aug 2023 - May 2025

Master's (M.S.) | Information Systems | **GPA: 3.86/4.0**

- Graduate Assistant (two departments); featured in CSUF [newsletter](#) for helping them develop and deploy TitanGPT.

A.P. Shah Institute Of Technology, Thane, Maharashtra

Aug 2019 – May 2023

Bachelor Of Computer Engineering (B.E) | **GPA : 8.39/10**

SKILLS & CERTIFICATIONS

ML & Gen Ai: LangChain, LangGraph, Prompt Engineering, HuggingFace, RAG, LSTM, Scikit-learn, PyTorch, TensorFlow

Data & Analytics: Power BI, Tableau, Pandas, NumPy, Plotly, SQL, Time Series, A/B Testing, Clustering, Alteryx

Data Platforms: Databricks, AWS (Bedrock, SageMaker), GCP(Vertex AI), Azure, PostgreSQL, MSSQL, Vector DBs

Tools & Collaboration: Git, GitHub, JIRA, VS Code, Notion, Agile, React, Team Communication, Confluence

Software & DevOps: Git, CI/CD, Docker, Kubernetes, MLFlow, Streamlit, JIRA, VS Code, React, RestAPI, FastAPI

Certifications: [Aws ML Specialty](#), [Microsoft Power BI DA](#), [Databricks DA](#), [Google DA](#), [CSPO](#), [Alteryx Core Designer](#)

WORK EXPERIENCE

Associate Consultant- AI/ML | KPMG | Mumbai, Maharashtra

July 2025 - Present

- Lead AI Architect for an enterprise-grade NLP grievance chatbot system for PFRDA on GCP Vertex AI, implementing custom intent classification and entity extraction to automate 20% of manual routing.
- Developed an automated NLP pipeline using Vertex AI to train and deploy a high-granularity classification model for a 60-class dataset of 40k+ records, optimizing feature preparation and model selection.
- Fine-tuned Gemini 1.5 Flash using SFT on a 60k-sample NLP dataset to automate a complex two-tier grievance classification architecture, significantly improving categorization granularity and downstream routing precision.
- Engineered a proof-of-concept Agentic AI system using LangGraph, designing multi-agent workflows that leverage MCP-based tool calling for deterministic financial decision-making to replicate *AlphaAgents* research paper.

Data Science Graduate Assistant | California State University | Fullerton, California

Aug 2024 – June 2025

- Developed a Retrieval-Augmented Generation (RAG) pipeline for TitanGPT, integrating internal research APIs to provide real-time external context and standardizing automated student FAQ resolution.
- Designed a LangGraph-based RAG workflow to systematically retrieve and synthesize research datasets, improving documentation and interpretability of model-driven insights.
- Built quantitative models on Vertex AI using large-scale datasets to identify default-risk drivers, incorporating model evaluation and backtesting techniques used in financial predictive modeling.
- Fine-tuned FinBERT to extract sentiment-driven financial risk signals from 10-K filings, similar to factor-based modeling used in quantitative market risk analysis.
- Helped automate XGBoost model training on AWS SageMaker, improving reproducibility and reducing setup time for student retention studies.

Analyst | The Esports Club | Bangalore, India

Oct 2020 - May 2021

- Performed customer segmentation by deploying k-means clustering to segment 100K+ survey responses, enabling a 30% boost in ad targeting efficiency for a client team.
- Executed A/B testing with statistical hypothesis tests to analyze the effectiveness of user engagement strategies, quantifying their impact on reducing churn risk by 20%.
- Created an interactive Tableau dashboard to analyze competitive player statistics across multiple-sports games, providing data-driven insights for player endorsements and sponsorship selection.

PROJECTS

Research Assistant – GenAI | Tools – [Github](#), [Python](#), LangGraph, OpenAI GPT-4o, React, FastAPI

- Built a LangGraph-powered multi-agent pipeline simulating OpenAI's Deep Research Mode for query resolution.
- Designed custom tools and prompts templates for coordinated analysis, fact-checking, and report generation.

DataViz Pro – AI | Tools – [Python](#), Streamlit, Plotly, LangChain, OpenAI GPT-4o

- Built a GenAI dashboard using LangChain agents and GPT-4o to answer CSV queries with dynamic Plotly visuals.

SmartCare AI – GenAI | Tools – [Python](#), Streamlit, LangGraph, OpenAI GPT-4o, MS SQL Server

- Optimized 120K-row joins in BigQuery & MS SQL for real-time doctor APIs, cutting cost 45% and latency 35%.
- Enabled real-time slot scheduling by mapping symptoms to specialists and querying availability via MS SQL Server.

Open Source Contributions | Tools – Langchain, CleanLab

- [LangChain](#) - Refactored LangChain's AWS Lambda tool to use AgentExecutor, improving support for LangChain v0.3.
- [CleanLab](#) - Enhanced Cleanlab's data-centric validation and **anomaly detection** for credit card fraud.