Prajwal Sathyanarayana

520-609-6114 | prajwals@arizona.edu | github.com/PRAJWAL2108 | linkedin.com/in/prajwals21 | prajwalsathyanarayana.github.io

Summary

AI Engineer with 3+ years of experience building and shipping AI/ML powered products end to end, including LLM-based systems and agentic workflows. Strong working understanding of prompt engineering, tool calling / APIs, retrieval / context construction, memory usage, and model evaluation.

Skills

- **Languages:** Python, SQL, C, Java, C++
- **Machine Learning & AI:** Large Language Models (LLMs), RAG, Vector Search, Embeddings, Prompt Engineering, Agentic Workflows, Model Evaluation, Tool Calling / APIs
- **Libraries & Frameworks:** Pandas, NumPy, Scikit-Learn, LightGBM, XGBoost, FastAPI, TensorFlow, PyTorch
- **Cloud & Developer Tools:** AWS (SageMaker, Lambda, RDS), Docker, Kubernetes, Terraform, CI/CD, GitHub Actions, PostgreSQL, REST APIs

Experience

Axio Education — JAN 2026 – Present

AI Engineer (Capstone Project INFO 698)

- Engineered an async FastAPI document intelligence pipeline using PyMuPDF and pdfplumber, achieving ≥88% table extraction accuracy and < 2s median processing time across 4 document modalities.
- Integrated Gemini 1.5 Vision multimodal LLM to generate anchored margin feedback with per-region confidence scoring, automating evaluation of 10+ questions per submission at ≥87% model confidence.
- Built a Gemini-powered document classification layer that auto-routes submissions to the correct grading pipeline, eliminating manual classification for 100% of self-contained assignment formats.

Banner Health — AUG 2025 – DEC 2025

Data Analyst (Externship with Eller Partnership Office, University of Arizona)

- Owned end-to-end design of AI-first workflows for complex knowledge work by building Python and SQL pipelines to extract, transform, and load encounter and procedure data, improving data accessibility by 35%.
- Designed reusable AI capabilities and evaluation pipelines by translating messy stakeholder needs into clear specs, prompts, and expected behaviors, enabling standardized reporting structures.
- Defined clear success metrics for AI behavior and automated reporting workflows using Python, reducing manual reconciliation effort by 40% while improving reliability of downstream surgical volume insights.

University of Arizona – ACT Lab — JAN 2025 – JUL 2025

Graduate Student Researcher

- Designed and iterated multi-step agentic workflows for reinforcement learning experiments, including planning, tool invocation, memory usage, and refinement loops to support large-scale simulations.
- Built LLM-adjacent data analysis pipelines in Python to debug, compare variants, and analyze real interaction logs from high-frequency simulations, improving experiment turnaround time by 30%.
- Implemented model output evaluations and monitoring strategies to detect common failure modes, maintaining accuracy and reproducibility across reinforcement learning pipelines.

Cognizant — JUL 2022 – JUL 2024

Data Scientist NOV 2023 – JUL 2024

- Addressed high data processing overhead across multi-source financial datasets by architecting scalable SQL and Python pipelines, reducing cloud compute and data preparation costs by 25%.

- Resolved costly false-positive default predictions by translating ambiguous requirements into strict engineering specs, improving detection accuracy by 18% and minimizing financial risk exposure.
- Targeted excessive AWS resource consumption from inefficient model inference by automating pipelines via SageMaker and APIs, cutting latency by 30% and lowering cloud infrastructure spend by 20%.

Jr. Data Scientist JUL 2022 – OCT 2023

- Tackled the high labor costs of manual revenue forecasting across business units by designing agent-like analytical workflows with tool-calling logic, reducing operational forecasting costs by 40%.
- Mitigated expensive administrative bottlenecks in tracking system metrics by developing automated Python dashboards, improving operational efficiency by 35% and heavily reducing reporting overhead.
- Prevented wasted operational spend on faulty A/B tests by implementing rigorous data quality checks and evaluation pipelines, stopping the costly deployment of underperforming variants.

Projects

Fare Estimation for Manhattan Yellow, Green Taxis — Aug 2025- Dec 2025

- Developed scalable data pipelines in PostgreSQL and Python to extract, transform, and load 3 years of NYC TLC data, enabling reliable fare prediction analytics.
- Automated data processing workflows and deployed a cloud-native FastAPI service, improving accessibility and usability of large datasets for downstream applications.

English Premier League Analysis — Jan 2025 - May 2025

- Simplified accessibility of complex datasets by developing interactive dashboards, increasing data clarity and usability for analysts and stakeholders by 40%.
- Automated reporting workflows integrating weather, odds, and match data, enabling consistent metric definitions across analytical views.

Predictive Modeling for Financial Markets — Aug 2024- Dec 2024

- Built structured data pipelines in Python to process historical market data, ensuring accuracy, consistency, and reliability across regression-based forecasting models.

Education

University of Arizona — AUG 2024 – MAY 2026

Master of Science in Data Science | GPA: 3.9

Dayananda Sagar University — AUG 2018 – JUL 2022

Bachelor of Engineering in Computer Science | GPA: 3.6