

ANISH NEHETE

Los Angeles, CA | anehete@usc.edu | LinkedIn | GitHub | Portfolio

EDUCATION

University of Southern California

Master of Science in Applied Data Science; GPA: 3.48/4

Los Angeles, CA

December 2026

Maharashtra Institute of Technology, WPU

Bachelor of Technology in Computer Science Engineering; GPA: 8.62/10

India

August 2024

SKILLS

Languages: Python, JavaScript, SQL, Java, C++ , C#

Backend & Cloud: FastAPI, Flask, React, Next.js, REST, WebSockets, AWS, SageMaker, GCP, Azure, Docker, Kubernetes, CI/CD, Git, pytest, Linux

Data & Databases: PostgreSQL, MongoDB, Redis, SQLite, Snowflake, Pandas, Pydantic, Spark, PySpark; familiar with Airflow, dbt, Kafka

ML/AI & Agents: PyTorch, TensorFlow, spaCy, medSpaCy, OpenCV, OpenAI Agents SDK, LangChain, CrewAI, AutoGen, MCP, MLflow

WORK EXPERIENCE

University of Southern California, Keck School of Medicine

Los Angeles, CA

Data Science Intern, Cancer Surveillance Program

September 2026 – Present

- Built an offline clinical NLP pipeline (Python, medSpaCy, HL7 v2) that parses HL7 pathology reports, strips embedded identifiers, and classifies cancer reportability under SEER/California rules, with every decision traced to its rule and triggering sentence.
- Deployed the classifier as an independent fourth assessor alongside CDC NOAH and existing registry checks in a casefinding workflow that screens approximately 240K pathology reports a year for LA County, flagging cross-system disagreements for registrar review.
- Engineered a PHI-safe dev workflow with a synthetic HL7 generator and pytest leak checks, validating the full pipeline with zero PHI exposure.
- Built a PySide6 case-review app rendering color-coded evidence (affirmed, negated, historical, excluded) and a step-by-step decision trace, cutting registrar verification time from minutes to seconds.
- Developed Power BI dashboards on pathology case volume, reportability trends, and invoice spend by funding source for registry leadership.

Pasadena Digital – Viterbi Startup Garage

Los Angeles, CA

AI/ML Engineer Intern

June 2025 – August 2025

- Optimized taxiway graph structures in OwlVision using OpenStreetMap, Osmnx, and NetworkX, applying cycle detection, centrality analysis, and pruning to eliminate 95% of redundant paths and improve geospatial clarity.
- Integrated a Visual Language Model into a PyQt5 annotation tool for prompt-based detection across 40+ aviation classes, cutting labeling by 60%.
- Fine-tuned Whisper-Medium on 12K+ hours of noisy ATC audio with hyperparameter tuning and RMS-based silence filtering, cutting WER from approximately 18% to under 8%; automated training and deployment with Roboflow and CI/CD.
- Built end-to-end CV and NLP pipelines (PyTorch, Dash, GeoPandas) for cockpit state detection and textual-cue parsing with decision support.

Plural Technology

Pune, India

Software Development Engineer Intern

January 2024 – July 2024

- Built and deployed an end-to-end React and FastAPI tool on AWS Amplify, scoping requirements directly with HR stakeholders and integrating Microsoft Teams and Power Automate to cut onboarding time by 50%.
- Automated visitor intake by deploying an AI virtual receptionist combining NLP, speech-to-text, text-to-speech, and facial recognition.
- Enhanced a transformer-based NLP model to recognize diverse accents via Python regex and custom entity rules, cutting processing time by 20%.
- Trained and deployed custom CV models (VGGFace, ResNet) on AWS SageMaker for facial recognition, decreasing inference latency by 20%.

ByteOrange

Pune, India

Data Science Intern

June 2023 – August 2023

- Automated business reporting with Azure ETL pipelines that staged cleaned data into Snowflake, powering auto-refreshing Power BI dashboards.

PROJECTS

The Sphere – Real-Time Geospatial Intelligence Platform | FastAPI, Next.js, PostgreSQL, Redis | Demo

March 2026 – April 2026

- Engineered a full-stack, near-real-time intelligence platform (FastAPI, Next.js) with modular microservices for retrieval, portfolio, and alerting behind typed REST and WebSocket interfaces.
- Designed a fault-isolated asynchronous ingestion pipeline polling 9 signal adapters via asyncio, normalizing and deduplicating events on 60-second cycles with per-adapter failure isolation.
- Implemented production-oriented persistence and resilience with PostgreSQL, Alembic migrations, Redis-backed rate limits and alerts, Docker Compose, startup health checks, and pytest, Vitest, and Playwright test suites.

Multi-Agent Trading System with MCP Orchestration | Python, OpenAI SDK, FastMCP, SQLite | GitHub

June 2025 – July 2025

- Orchestrated 4 concurrent trader agents with per-agent researcher sub-agents (up to 8 LLM agents per cycle) using the OpenAI Agents SDK with configurable GPT, DeepSeek, Gemini, and Grok.
- Integrated 6 MCP servers over stdio for portfolio operations, market data, and web research, so agents execute validated trades and log rationales.

ChatDB – Natural-Language Interface for SQL & NoSQL | React, Flask, Mistral 7B, MySQL, MongoDB | GitHub

January 2025 – May 2025

- Led a 3-person team building a React and Flask assistant that turns plain-English questions into schema-grounded MySQL and MongoDB queries.

Yelp Hybrid Recommendation System | PySpark, XGBoost, Matrix Factorization, NumPy | GitHub

August 2025 – December 2025

- Built an RDD-only PySpark recommender combining item-item collaborative filtering, 80-factor biased matrix factorization, and XGBoost in a ridge-learned ensemble, reaching approximately 0.98 validation RMSE.

PUBLICATIONS AND PRESENTATIONS

Unity ML-Agents: Revolutionizing Gaming through Reinforcement Learning | IEEE WCONF, 2024

- Implemented PPO and SAC-based RL models, tuning reward functions and policy/loss dynamics to train adaptive NPC behaviors.

Comparative Analysis of Object Identification Techniques for Bird Species | IEEE ASIACON, 2024

- Built a detection pipeline via YOLOv5 and YOLOv8 to identify rare species, improving live biodiversity monitoring.

Deep Learning for Exoplanet Exploration: Detecting New Worlds and Evaluating Habitability | ICDAI, 2024

- Built ANN and XGBoost models on NASA exoplanet data to detect exoplanets from light-curve patterns and predict habitability.

CERTIFICATIONS

Microsoft Certified Associate: AI Engineer (AI-203), Data Scientist (DP-100), Data Engineer (DP-203), Power BI (PL-300) | Transcript