

TIRTH ALKESH PATEL

tirth.patel.2026@anderson.ucla.edu | +1 972-900-4650 | Los Angeles

[LinkedIn](#) | [Portfolio](#) | [GitHub](#)

EDUCATION

UCLA ANDERSON SCHOOL OF MANAGEMENT

Los Angeles, CA

Master of Science in Business Analytics (MSBA)

Expected Dec 2026

Machine Learning for Decision Making, Statistical Foundations of Analytics, Operations Analytics, Data Engineering, Competitive Analytics

INDIAN INSTITUTE OF TECHNOLOGY, MADRAS

Chennai, India

Bachelor of Technology, Chemical Engineering

May 2025

Mathematical Foundations of Data Science, Molecular Data Science, Advanced Operations Research, Operations Management, Economics

SKILLS

Languages: Python, SQL, R

AI & ML: PyTorch, TensorFlow, CUDA, scikit-learn, XGBoost, Deep Reinforcement Learning, NLP (VADER/NLTK), LangGraph, LangChain, ChromaDB

Data Engineering: Apache Spark (PySpark), Apache Airflow, Snowflake, GCS, BigQuery, ETL Pipelines, dbt

Analytics & Viz: Tableau (incl. Geospatial Maps), Pandas, NumPy, Matplotlib, Power BI, MS Excel

Tools: Git/GitHub, MySQL, Docker, FastAPI, Claude Code, VS Code

PROFESSIONAL EXPERIENCE

Seat of Joy, India (Stealth Startup)

June 2023 – Feb 2024

Product Analyst (Aug 2023 – Feb 2024)

- **Market Sizing Model:** Built a custom statistical framework from scratch using 100+ Census of India 2011 tables (200K+ rows each) — combined birth-order probability distributions with linear regression to predict future customer cohorts year-by-year, improving market estimates by 45%; became the analytical backbone of all investor pitches and GTM planning.
- Designed an operations research optimization model for state-wise product distribution using elastic-net-style penalty terms to prevent market over-concentration — translated OR methods into a practical startup decision tool; delivered distribution recommendations grounded in the custom demand model rather than intuition.
- Gathered and analyzed behavioral data from 50+ customer interviews — synthesized qualitative findings into structured datasets that surfaced insights challenging the team's assumptions on pricing and positioning, directly reshaping the product strategy.
- **Shark Tank India Audition:** Developed and delivered a comprehensive data-driven presentation integrating market analysis, competitive benchmarks, and statistical projections; communicated complex analytical findings clearly to executive judges in plain, accessible language — advanced through 2 selection rounds out of hundreds of applicants.
- Performed competitive data gathering and analysis across 10+ market players; cleaned and compiled findings into structured reports that surfaced actionable strategic insights for go-to-market planning.

PROJECTS

Course RAG Pipeline — Agentic Document Intelligence System — *LangGraph, ChromaDB, FastAPI, LangChain, Docker, OpenAI API, Claude API*

- Designed and deployed a production agentic RAG system — an AI tool actively used by the UCLA MSBA cohort — to speed context discovery across 4 concurrent courses of lecture slides, transcripts, and PDFs; specified what each of the system's 13 LangGraph nodes should produce across 5 distinct query-routing branches and verified output correctness at every processing stage.
- Engineered a self-verifying deadline extraction pipeline that re-queries ChromaDB with a rephrased search after initial LLM extraction, cross-references both result sets to highlight date incongruencies, and surfaces conflicts with a confidence indicator — directly addressing the query type where a wrong answer is more harmful than no answer; implemented dual-provider LLM fallback (Claude primary, GPT-4o-mini secondary) to ensure system reliability for real users.
- Deployed the full stack (FastAPI + WebSocket backend, self-hosted ChromaDB, Google Drive integration, vanilla JS frontend) on Oracle Cloud; resolved 9 distinct production failures across Linux networking, Docker credential management, and API timeouts; built a human-in-the-loop upload approval flow enabling users to verify AI file organization decisions before execution — communicating AI reasoning in plain language for non-technical users.

Maltese Gear Cube Solver — Deep Reinforcement Learning — *PyTorch, CUDA, NumPy*

- Applied the DeepCubeA (Deep Approximate Value Iteration) algorithm to a custom Maltese Gear Cube — a domain it was never designed for — building the full system from scratch with a 144-element state representation and 816-dimensional one-hot encoding; adapted the algorithm's core assumptions to fit a state space with fundamentally different structure from standard benchmark puzzles.
- Migrated the training pipeline from CPU/NumPy to a fully GPU-resident PyTorch/CUDA implementation on RTX 5070 Ti for significant throughput gains; engineered Weighted A* search guided by a learned neural heuristic using the two-network DAVI algorithm, replacing hand-tuned rules with an AI-driven approach and benchmarking solution quality across parameter configurations to validate reliability.

Weather & Dining Intelligence Pipeline — *PySpark, Snowflake, Airflow, GCS, Tableau, VADER NLP*

- Built an end-to-end data pipeline fusing 2M+ Yelp consumer reviews with NOAA weather records across 9 US cities to investigate how external conditions influence dining sentiment — applied VADER NLP for behavioral pattern detection and built a full Spark-based medallion pipeline (Bronze → Silver → Gold) with Airflow orchestration to make the analysis reproducible and scalable.
- Designed a Snowflake star schema for cross-dataset querying; built an interactive Tableau dashboard with KPI tiles and geospatial visualizations to surface behavioral patterns as clear, actionable insights for stakeholders.