

+1(619)9538035

San Diego, CA

nmahor@ucsd.edu

emailfornishchay@gmail.com

Nishchay Mahor

Data Scientist/ML Engineer

contextjetai.com/nishchay

medium.com/@emailfornishchay

github.com/NishchayMahor

linkedin.com/in/nishchaymahor

EDUCATION

Master of Science (M.S) in Data Science: GPA: 4, *University of California, San Diego*

July 2027 (Expected)

Bachelor of Engineering (B.E) in Computer Science: GPA: 3.9, *Chitkara University*

July 2023

TECHNICAL SKILLS

Languages: Python, SQL, C++

GenAI / LLM: OpenAI/Azure OpenAI, LangChain/LangGraph, RAG, Tool-Calling, Agentic Flows, ML Systems

ML / AI: scikit-learn, CatBoost/XGBoost, time series, pytorch, keras, MLflow

Platforms & Tools: Git, Databricks, AWS, Azure, Terraform, Neo4j, Next.js, LaTeX

EXPERIENCE

Data Scientist

June 2024 - July 2025

CSi.ai (previously Course5 Intelligence)

Bengaluru, IN

- Led a cross functional team of engineers and analysts under the **Centre of Excellence** team to build enterprise AI products for **Marketing, CPG, and Tech** use cases, contributing **\$500K+** in project revenue in FY 2024 - 2025.
- Built **Multimodal Synthetic Market Surveys** - a framework to synthesize real-time respondent data for text and image studies with ~90% accuracy vs live benchmarks (MAPE < 10% across Likert items); generated >**\$300,000** in revenue; built with **LangChain, Azure OpenAI, multi-agent pipelines, and Apify**; deployed on **Azure Web App** via **GitHub Actions**.
- Developed an LLM Based **Virtual try-on (VTON) Assistant** for a leading clothing brand that drove **time on page +25%, CTR +18%**; built using diffusion-based image try-on (**StableVITON from Hugging Face Hub**) and LLM product/size chat using **OpenAI(gpt-image-1), LangGraph**; CV pipeline with **MediaPipe** and RAG over catalog (**Pinecone**);
- Reduced deployment load >**80%** by building CI/CD pipelines for POC apps; automated releases to **Azure Web Apps** and cut **time-to-demo ~50%**, accelerating client sign-offs; stack: **GitHub Actions, Docker**; multi-cloud: **Azure, AWS**.
- Created and maintained an **in-house LLM stack**(repos, templates, caching, prompt libraries); enabled rapid development of **5+ POCs**, sped up model development by **40%**, and **reduced over 60% in API Costs**.

Data Analyst

May 2022 - June 2024

CSi.ai (previously Course5 Intelligence)

Bengaluru, IN

- Last-mile delivery (Supply Chain):** built a **Neo4j/Cypher knowledge graph** for routes, hubs, SLAs; **on-time delivery +8%, route time -12%**, faster exception resolution; APIs exposed via **FastAPI**.
- PAM + planning (Industry 4.0):** anomaly detection on vibration/temperature for conveyors + **RAG + forecasting** for planners; **unplanned maintenance -15%, planning cycle -30%**; stack: **scikit-learn/Prophet, LangGraph, Pinecone, Databricks**.
- CLTV (OEMs in semiconductor ecosystem):** **BG/NBD + Gamma-Gamma** models for laptop/cloud channels; **MAE -18%** on spend prediction and **campaign ROI +15%**; productionized on **Databricks + MLflow**.
- Search-rank prediction (major tech client):** simulated marketplace ranking signals (keywords, price, reviews sentiment) and built **XGBoost/CatBoost** models; **~70%** daily rank prediction accuracy → higher **Share of Search/Shelf** across US/CA retailers.

Co-Founder

March 2025 - September 2025

ContextJet.ai

Bengaluru, IN

- AI Sales Development Representative:** automated prospecting, enrichment, personalization, outreach, and reply handling for a Wall-Street-focused brand targeting PE/hedge funds/family offices; **qualified meetings +35%, manual prospecting -60%, 200-300 targeted leads/week**. Built with **LangChain + Pydantic agents, Apify** for scraping, **React UI**; **deployed on Azure App Services**.
- Knowledge GraphRAG platform:** entity-linked graph over docs for import/export workflows; **answer precision +87%, research time -45%**, auditable citations and faster handoffs; built with **LangGraph, vector DB, Python, Salesforce integration**; deployed on **Azure**.
- Megan (AI IVR + tele-assist):** IVR with real-time ai video consult and auto-logging for pre/post-procedure guidance; **call deflection ~45%, 3-5 hrs/week/provider saved**; built with **Retell/Twilio, FastAPI**; deployed on **AWS (API Gateway + Lambda)**.

PUBLICATIONS

- Keynote Speaker/Author "**Revolutionizing Market Surveys through Unprecedented Generative AI for Efficient Data Synthesis**" presented at the Machine Learning Developers Summit 2024 and published in **Lattice Journal (AoS), Vol. 5 Issue 1, Jan 2024**.
- Author/co-author "**From Risk to Safety: AI's Promise for India's Road Safety**" published in **IJRDO - Journal of Applied Management Science, Jan '24, Vol 10, Issue 1**.