DIPJYOTI DAS

Charlotte, NC, USA • 651 747 6515 • dipjyoti@gmail.com
https://www.linkedin.com/in/dipjyotidas, https://github.com/dipjyotidas

Professional Highlights

● **13+** years of hands-on Experience, delivering end-to-end **Gen AI/Machine Learning** & Data Science solutions; strong storytelling skills with a track record across startups & Fortune 150 companies - Tech, Finance, Utilities, Healthcare, Insurance, Telecom, Logistics.

● Delivered projects; managed clients across cross functional business units – Marketing, Sales, Product, Operations, Engineering etc.

● Mentored a cohort of 20 professionals in the PG AI/ML program on Great Learning's EdTech platform in partnership with UT Austin.

● Speaker & session Chair at AIM 2026 SFO, Ai4 and REWORK AI conferences on Production-scale Agentic AI systems in Enterprise (YouTube: AIM 2026: https://www.youtube.com/watch?v=zOum0De_6K0 & Ai4: https://www.youtube.com/watch?v=K_5Ij0OawjI).

● Led technical projects; Managed & mentored teams of 1–6 across multiple initiatives. Acted as a technical advisor to AI engineers & data scientists and served as a thought partner to executives balancing hands-on delivery with leadership responsibilities.

● Hands on AI/Machine & Deep Learning, **Generative AI** model development (Python,Tensorflow, PyTorch) & deployment (K8s): built Time Series Forecasting, Chatbot, Propensity, Clustering, Recommendation models with AWS Sagemaker, Bedrock; Azure AI Foundry

● Deployed Propensity models to SAP cloud with architects; automated ETL and data engineering pipelines using PySpark/Spark SQL. Production model generated $240M in incremental annual revenue at Brighthouse Financial.

Skills: Python (keras, scikit-learn, networkx), **AWS** (Jumpstart, Bedrock, Browser use, Computer use, Lambda, OpenSearch, EKS, RDS), **Azure** AI Foundry, Azure Function, App Insights, Databricks**,** GCP Vertex AI, **Kubernetes**, Docker, Microservices, Snowflake, PostgreSQL, Argo, Azure DevOps, Github actions, CircleCI, Elastic Search, FastAPI, Flask, Pyspark, Streamlit, Tensorflow, Hadoop

Machine/Deep Learning/Model serving: Regression (Lasso, Ridge), Classification (Logistic regression, Random Forest, KNN, Naive bayes), Boosting (XGboost, LightGBM), Ensemble, Clustering (Gaussian mixture, DBSCAN, K-means), Graph Network, Neural Network (RNN, LSTM, CNN, Encoder-Decoder), Diffusion models, **MLOps**, CI/CD/CT, **Databricks** Model registry, MLflow, TeamCity, ArgoCD

Generative AI: Generative Adversarial Networks, **Transformer**, BERT, LLM - GPT-4o, Gemini, Llama 3, Mistral, Claude 3, FinGPT, LangChain, Llamaindex, LangGraph, Google Agent Dev Kit, Agentspace, Crew AI, FastMCP, A2A protocol, **Langfuse**, Embeddings – Hugging face, Cohere; Vector DB – Vespa, OpenSearch; Guardrails, Ragas, GEval, OpenAI Moderation API, NLP, RLlib, Gymnasium

Certifications: AWS Machine Learning Specialty & AWS Solutions Architect Associate (in progress)

Work Experience

1. Gen AI Architect, Ascendion, Charlotte, NC — Oct 2025 - Present

Ascendion is a software engineering firm delivering enterprise Generative and Agentic AI solutions. Led architecture and production deployment of two large-scale AI systems across Telecom and Healthcare enterprises.

● Led cross-functional engineering teams to Architect and Deploy a unified **Agentic AI** chatbot microservice on AWS EKS within **Charter Communications'** Tech Mobile application, serving 80,000+ field technicians, supervisors, maintenance engineers.

● Architected LangGraph-based orchestration to route role and intent aware queries across task-specific agents (Job Insight, Tech Assist, APIs) with Langfuse observability; reduced support calls by 60%, accelerated service completion with million-dollar savings.

● Architected stateless REST services with FastAPI, OAuth/JWT security, and PostgreSQL/Redis-backed state management to support scalable enterprise workloads.

● Design and lead an Agentic AI solution (LangGraph) for **CareSource** to automate Discharge logs fax intake using Azure Document Intelligence and LLM validation (HITL), integrating eligibility and Prior Authorization APIs to reduce manual review; improve auditability.

● Led serverless deployment with Azure Functions & CI/CD; implemented observability via Application Insights and centralized prompt versioning, logging. Oversaw React-based dashboard deployed on Azure Web App to enable business review & HITL validation.

● Owned end-to-end solution architecture, client demos, serverless and AWS EKS deployment in Dev/Prod clusters, CI/CD pipeline for microservices and infrastructure management; perform code reviews, guide technical design, and lead hands-on POCs.

● Collaborate closely with clients on delivery management, **roadmap** planning, KPI tracking, business use case definition; advise stakeholders on build/buy/rent strategies, LLM and embedding model selection, vector database & observability tools.

● Led development of a **RAG** service by extracting and embedding 1,000+ confluence pages and 5,000+ technical attachments (images, pdfs, excel, word) into PGVector DB, preserving document hierarchy and multi-level contextual relationships.

DIPJYOTI DAS

Charlotte, NC, USA • 651 747 6515 • dipjyoti@gmail.com
https://www.linkedin.com/in/dipjyotidas, https://github.com/dipjyotidas

2. AI Architect/Engineer, CGI Technologies, Charlotte, NC June 2023 – Oct 2025

Worked in CGI's US IP core AI product team- Pulse AI https://www.cgi.com/en/solution/cgi-pulseai. GenAI services - built with open-source solutions, self hosted as scalable microservices in **Azure Kubernetes Service** (a) Ask Data - Advanced RAG integration of **Onyx** -Danswer AI (b) Ask Insight (text to sql) – integration of Vanna AI (c) GenAI inference - self hosted Llama/Mixtral with vllm and Ray serve (d) Monitoring (Prometheus, Grafana, ELK) (e) Responsible AI- Prompt guard

● **CGI Agentic AI** platform**:** Built & optimized LangGraph React agents/tools using design patterns (Reflection, Planning, Multi-Agent) for tasks: web search (Serper), productivity (Jira, Office 365), database (Postgres), utilities (REPL, Arxiv, Wikipedia, custom functions).
● Developed Fast APIs for Agent AI config console (no-code UI) and observability using **Langfuse**; deployed via Helm in Kubernetes. Developed default and custom evaluation metrics (Ragas, G-Eval, Json) by integrating Langfuse APIs with prompt versioning.
● Served as **AI Architect** for **Fannie Mae**'s Enterprise Gen AI team (AWS Bedrock/Sagemaker); defined GenAI **roadmap**, solution vision, milestones across business units. Built reusable accelerators, production-ready AI solutions like chatbots, knowledge assistants.
● Designed enterprise GenAI architecture in AWS with LLM Model Garden layers; recommended LLMs (Llama 3, Mistral, Claude Sonnet), embedding models (Cohere, Titan), vector DBs (OpenSearch, Milvus), benchmark with MMLU, Chatbot Arena, HumanEval
● Built Agentic AI system to automate Foreclosed Property Valuation, integrating appraisal products, Collateral Underwriter, Zillow - reducing valuation approval time by 70% through multi-agent orchestration for document ingestion, validation, red-flag detection.
● Developed multi-agent workflows within **AT&T**'s Ask ATT platform for customer service automation, including a Team Planner Agent orchestrating sub-agents (Ask Docs API, Ask ATT, Ask Web) with task-specific prompts; Validator Agent for output quality/compliance.

3. Staff Data Scientist, One Concern**,** Menlo Park, CA Nov 2021 – May 2023

One Concern is a Data & Insights Risk Analytics Climate Tech startup. As part of Research & Solutions team, we worked with external clients (Insurance, Banking, CRE), Eng, Go to Market, Sales team to build POC's for clients & productize/scale data science solutions.

● Delivered Resilience & Business Interruption scores for U.S./Japan commercial properties via REST API. These analytics powered 1C-DNA and 1CRX products later integrated into Swiss RE CatNet and Willis Tower Watson platforms through strategic partnerships.
● Data & Analytics product: Develop Exceedance probability/Downtime statistics production pipeline across hazards, return periods, design wireframe; scaled Resilience metrics for 30M properties with Eng team using Kubernetes (GCP), Argo, and CircleCI (CI/CD).
● Mentored data scientists and interns/new hires, led code reviews, and supported pre/post sales client engagements (used Tableau).

4. Data Scientist II, Duke Energy, Charlotte, NC Mar 2019 - Nov 2021

● Built Word2Vec model for failure mode synonyms, used Elastic Search to query work orders, forecast failure mode, detect anomalies
● Provided technical leadership, mentorship to junior data scientists on solutions for water heater failure prediction & sales conversion.
● Developed short/long-term Forecasting models (LSTM, CNN1D, SARIMA) for daily/weekly Disconnect Non Pay work orders using features like seasonality, weather, bill cycles; deployed via batch cron jobs on windows server using **CA Workload** Automation.
● Developed & deployed a batch Risk Propensity ML model to score customers by likelihood of delinquency. *Est Impact: $5.1M/year*
● Developed an unsupervised Gaussian Mixture Model to recommend and rank commercial properties for Energy Efficiency programs in non-native territories; managed and led a team of two data scientists. *Est financial impact: $8–10M/year.*

5. Data Scientist II, Brighthouse Financial, Charlotte, NC June 2017 - Jan 2019

● Built Propensity model from diverse data sources to rank & score Financial Advisors most likely to sell Flex & Shield annuity products.
● Built a 3-layer stacked ensemble model (Logistic Regression, XGBoost) across advisor segments and product models; evaluated email campaigns via A/B testing; collaborated with Model Risk/Compliance teams. *Impact*: *$60M in incremental quarterly sales revenue*.
● Led the migration of ML models to SAP cloud in collaboration with architects; deployed Propensity models to dev/prod clusters, automated production workflows (Python), built ML & ETL pipelines (used PySpark), and onboarded the team to the SAP cloud env.

6. Decision Specialist, R+L Carriers, Ocala, FL March 2013 - June 2017

● Built Time Series Forecasting models (Holt Winter, ETS, ARIMA) using VBA, R to predict key metrics- revenue, shipments, etc
● Automated Forecasting models using SQL/R, reducing run time to 1 hour and saving *$800K annually* by eliminating manual effort.

Education

● **Master of Science**, Materials Science and Engineering, University of Florida, December 2012. GPA – **3.9/4**
Courses: Machine Learning (Stanford); Error Analysis, Optimization & Experimental Design; Data Science: Data to Insights (MIT)
● **Bachelor of Technology**, Metallurgical & Materials Engineering, National Institute of Technology, Trichy, India 2010. GPA - **8.3/10**

Publications: "Effects of cycling on the pseudoelastic properties of CuAlMnNi and TiNi based pseudoelastic alloys", IJSCS, Vol 1, 2009
https://oaktrust.library.tamu.edu/server/api/core/bitstreams/6c7c7c05-6407-4351-88aa-4b398794014c/content