Himanshu Laddhad

765-414-2728 | himanshuladdhad11@gmail.com | linkedin.com/in/himanshu-laddhad | github.com/Himanshu-Laddhad

SUMMARY

Data Scientist with 3+ years in enterprise analytics and AI engineering — causal ML, model fine-tuning, and LLM orchestration in production.

Certifications: Tableau Desktop Specialist, Microsoft AI-900, AWS CLF-C02, DataCamp: Data Scientist Associate & SQL Associate

Languages & Libraries: Python (Pandas, NumPy, Scikit-learn, XGBoost, LightGBM, HuggingFace Transformers, spaCy, PySpark), SQL (PostgreSQL, BigQuery), DuckDB, XSLT, XML, EDI, Parquet, PyTorch, AWS Sagemaker, ML (Supervised, Unsupervised), ETL Pipelines

AI, MLOps & Platforms: LangGraph, LLM APIs, RAG, QLoRA/PEFT, ChromaDB, LangChain, LlamaIndex, MLflow, Weights & Biases, SAP Datasphere, SAP BW, dbt, Power Automate, Azure Data Lake Gen2, Databricks, Snowflake, Git

Methods & Domain: Causal Inference & Uplift Modeling, A/B Testing & Experimentation, Dimensional Modeling, Data Contracts, Feature Engineering, NLP & Transformers, SHAP & Interpretability, Root Cause Analysis, Data Quality Management, Agile/Scrum

EDUCATION

Purdue University, Daniels School of Business — West Lafayette, IN
Master of Science in Business Analytics and Information Management, GPA: 3.7/4.0 — August 2025 – July 2026

Sardar Vallabhbhai National Institute of Technology — Surat, India
Bachelor of Technology, Mechanical Engineering, GPA: 3.6/4.0 — August 2019 – May 2023

PROFESSIONAL EXPERIENCE

PartnerLinQ — West Lafayette, IN
AI Engineer - Industry Practicum — January 2026 – May 2026

- Architected a RAG-grounded EDI-to-XML mapping agent routing messages through Groq to a GPT-4.1 tool-calling model, retrieving from a **400+ file** XSLT corpus keyed by company and function - cutting mapping build time from **3-5 weeks** of manual development to **3-4 hours**.
- Trained an SLM (Qwen) with QLoRA covering transform generation, modification, repair, and explanation, grounded by the same retrieval layer at **73%** first-pass Altova validity against the GPT-4.1 pilot's **78%** - removing per-call API dependency across the full task set.
- Delivered the Streamlit evaluation harness fronting both pipelines, executing generated XSLT against user-supplied test cases in-session - handed off as the reference implementation for PartnerLinQ's software engineering team to productionize.

JCB India Ltd — Pune, India
Data Scientist - Business Intelligence and Quality — February 2024 – July 2025

- Owned the curated analytics layer on **SAP Datasphere**, modeling four event streams from Parquet on ADLS Gen2 into a conformed star schema across INBOUND → HARMONIZED → CURATED layers at **35K rows/day** and **670-700 attributes** per entity on a single serial join key - replacing manual consolidation for **10+** cross-functional consumers.
- Engineered the **dbt** build and test suite on the dbt-hana adapter, enforcing uniqueness, referential, and accepted-value contracts in SQL plus custom tests for stage-sequence and N-sigma drift - quarantining failed rows with typed exception codes.
- Designed the **Power BI** consumption layer over SAP HANA ODBC on hourly incremental refresh partitioned on ingest timestamp to absorb late-arriving corrections, surfacing **15+ business KPIs** across 3 stakeholder groups - replacing **20+ hours/week** of hand-built reporting.
- Orchestrated severity-tiered exception alerting through **Power Automate** to Teams adaptive cards with writeback acknowledgement and per-rule cooldowns - cutting exception rates **15%** against an open-vs-acknowledged SLA view.

JCB India Ltd — Pune, India
Data Scientist - Supply Chain — May 2023 – February 2024

- Led a **3-person** rebuild of plant-level demand forecasting from Excel to **XGBoost** in Python on SAP BW, engineering lag, volatility, capacity-headroom, and pipeline-weighted features across **500+ SKUs** and 4 product lines - cutting MAPE **18%** against the spreadsheet baseline.
- Devised a two-track strategy in Python routing regular-demand SKUs to the model and intermittent-demand SKUs to a trailing annual average, validated on a temporal split with a 6-month holdout - replacing gut-based planning for procurement and capacity teams.
- Built a Python BFS traversal over a multi-level component hierarchy, exploding SKU forecasts into per-component monthly demand consumed directly by procurement - holding **98% on-time delivery** across **500+ SKUs** vs a 95% benchmark.
- Deployed Python time-series monitoring flagging demand deviations **4-6 weeks** ahead across **6+ international markets**, with automated alert routing into operational dashboards - replacing reactive post-hoc review for planning teams.
- Instituted freshness validation and output sanity checks as the Python pipeline's first operations after a stale input produced silently wrong forecasts - halting bad data before it reached procurement.

Bosch — Nashik, India
Data Analyst (Operational Analytics) — April 2022 – August 2022

- Structured a Pareto analysis over 9 months of operational time-series data in Python, ranking **140+ exception categories** by contribution to downtime - lifting process efficiency **10%** for the operations team.
- Integrated corrective-action tracking against the ranked categories across procurement, quality assurance, and operations - cutting exception frequency **~90%** and reprocessing rate **14%** on a high-volume workflow.

PROJECTS

- **Causal Promotion Intelligence System** *(Python, EconML, DuckDB, Scikit-learn) Built an uplift modeling system estimating heterogeneous treatment effects via CATE meta-learners on RCT data - segmenting customers by promotion sensitivity for causal campaign attribution.*
- **Customer Churn & Lifetime Value Intelligence** *(Python, XGBoost, LightGBM, DuckDB, SHAP, Optuna) Engineered 35 features from* **410M log rows** *across* **6.7M users** *in DuckDB SQL; XGBoost* **ROC-AUC 0.9815**, *LightGBM CLV* **R^2=0.9923** *- EV model projected* **$1.26M ROI**.
- **Retail Review Intelligence - Aspect Sentiment Analysis** *(Python, DistilBERT, HuggingFace, spaCy, LIME) Fine-tuned DistilBERT on* **50K** *reviews for 3-class sentiment against VADER, TextBlob, and TF-IDF baselines; aspect-level sentiment via spaCy with LIME attribution.*

AWARDS AND AFFILIATIONS

- JCB - Excellence in Planning and Efficiency (2025), Rising New Joinee Award (2024)
- Indian Institute of Management Ranchi (Tier 1 Business School, India) - Research Intern: Market & Competitive Analysis of EV Landscape
- Purdue University, Daniels School of Business - Graduate Teaching Assistant, AI for Business (June 2026 - July 2026)