

CAREER SUMMARY

Senior Data Scientist with 6+ years applying scientific methods to high-stakes business problems, building production-grade ML/DL and end-to-end LLM/Agentic AI systems at scale. Hands-on across time-series forecasting, causal inference, recommendation, and large-scale data, with a product-oriented mindset and a habit of turning research into real-world impact.

WORK EXPERIENCE

Sr. Data Scientist – Aviso AI, India

May '25 - Present

- **Multi-Agent LLM System for Sales Intelligence & Recommendation**
 - Architected a **multi-agent LLM system** processing **100K+ deals daily** across 15 enterprise clients, replacing rule-based logic with insight generation from CRM and engagement data, governed by evaluation and safety guardrails
 - Replaced correlation-based explainability with a causal inference backbone using **Double Machine Learning** and **Causal Forests** to surface actionable recommendations, driving a **~12% increase** in sales representative engagement
- **Built and scaled the LLM engineering function**
Defined the technical hiring bar and **grew the team from 1 to 5 engineers** to support end-to-end development and production deployment

Sr. Computer Vision Engineer – Bosch Global Software Technologies, India

August '22 - April '25

- **Advanced Vision Model Research & Publications**
Published a research paper and filed a patent on enhancing foundational models (EVA-02), leveraging 90%+ unlabeled Bosch data to **boost performance by ~5%**. Collaborated with Bosch Romania to further **improve EVA by ~4.31%** via cluster analysis and hypothesis testing on **DINOv2** embeddings
- **End-to-End Vision System Development**
Built a Vision Assignment Marketplace, **training models like SWin, DeiT, and DeTr** with **~6.72% average performance deviation**. Designed an image segmentation pipeline using ConvNeXt on fisheye images, **achieving an F1-score of ~57%**
- **Real-Time Edge Deployment & Codebase Optimization**
Deployed an **object detector and motion analyzer** on a Jetson Nano edge test-bed with **~30 FPS** latency. Refactored legacy code to support newer vision transformer architectures, **improving adaptability and efficiency by 14%**

Data Scientist II – HiLabs Inc., India

October '20 - July '22

- **Scalable Data Pipeline & Anomaly Mitigation for Anthem**
Engineered a multi-million dollar data pipeline to transfer anomalies from large-scale SQL/RDBMS sources to HBase and Solr for remediation via a ticketing system, **improving Data Quality Index (DQI) from ~78 to ~86**. Refactored anomaly mitigation process, projected to **prevent over \$200M in annual losses** for a leading U.S. health insurer
- **Production-Scale OCR & GAN-Based Contract Processing**
Deployed an enterprise-grade OCR solution **using NLP** and custom algorithms on USP contracts after enhancing image quality via GANs, **generating over \$100K in business value**

Machine Learning Engineer Intern – GMO Research, Japan

Jan '20 - March '20

- **Survey Targeting optimization with Time-Aware ML**
Forecasted 6-hour survey response likelihood using time-aware Random Forests, Decision Trees, and DNNs. Addressed class imbalance via SMOTE and custom metrics, **achieving 91% efficacy** on responsive marketing panels
-

SKILLS

- **Programming Languages and Libraries:** Python, SQL, Scala, Pandas, Numpy, Scikit-Learn
 - **Big Data & MLOps:** Spark, Hadoop, HBase, Solr, AWS, Docker, Git, CI/CD, GPU/HPC
 - **ML/DL/LLM Skills:** PyTorch, TensorFlow, Transformers, LangChain, LangGraph, RAG, Agentic AI, MCP
 - **Models and Paradigms:** Time-Series Forecasting, Recommendation, Causal Inference, Anomaly Detection, Computer Vision
-

PROJECTS

- **Deep-Learning Architectures from Scratch** – open-source PyTorch library implementing 25+ computer-vision architectures (CNNs like ResNet and ConvNeXt, Vision Transformers like DeiT and Swin, and segmentation networks) from their original papers; modular, packaged codebase with unit tests, CI/CD via GitHub Actions, config-driven training, experiment logging and metrics, and hosted [documentation](#)
-

PUBLICATIONS

- **Co-authored a patent leveraging dimensionality reduction techniques** to detect and localize failure regions in segmentation outputs from foundation models, enhancing model interpretability and reliability
 - **Published research demonstrating a two-stage fine-tuning approach** that improves foundation model performance on segmentation tasks by over **5% mIoU** using unlabeled data, advancing semi-supervised learning strategies
-

EDUCATION

Indian Institute of Technology, Mandi

August '16 – July '20

B.Tech in Mechanical Engineering, Minor in Intelligent Systems