

Hargurjeet Singh Ganger

gurjeet333@gmail.com | +91 9035828125 | Bangalore, India
[LinkedIn](#) | [Github](#) | [Medium](#) | [Portfolio](#)

Summary

Experienced IT professional with **15+ years** in the industry, specializing in **data science, statistical analysis, and machine learning and Generative AI**. Expert in LLMs, AI model development. Proficient in **Python, SQL**, and cloud platforms like **AWS** and **GCP**, with expertise in building and deploying scalable ML and Agentic solutions.

Skills

Statistics & Machine Learning– Regression analysis, statistical inference, hypothesis testing, supervised and unsupervised learning, ensemble methods, feature engineering, boosting trees, model evaluations.

Model Evaluation & Interpretability: SHAP, LIME, Explainable AI, Precision-Recall, ROC-AUC, RMSE, MSE,, Prompt Alignment, Answer Relevancy, Faithfulness, Context Precision/Recall, Hallucination, Bias, Toxicity.

Generative AI - Vector DB, RAG patterns, Guardrails, PII Filtering, Agents, Fine-tuning, Observability, MCP

LLM Framework and API's - Langchain, Langgraph, CrewAI - tool augmented workflow, multi agent systems.

MLOps - Model deployments, Monitoring and version control, CI/CD pipelines, Feature store, Data drift, Model Governance, Docker, FastAPI, Gitlab, MLflow, YAML, Alerting and Logging, OpenAI, Anthropic, Gemini API's.

Cloud Platforms - AWS (SageMaker, Lambda, Bedrock, Textract, OpenSearch, Step Function, Cloud Watch)

NLP - Dense/Sparse Embeddings, Named Entity Recognition(NER), Topic Modelling, BERT/DistilBERT.

Programming & Tools: Proficient in **Python** (pandas, NumPy, Scikit-learn, TensorFlow, PyTorch), **SQL, PySpark**, and **SparkSQL, Linux** for end-to-end data workflows and machine learning pipelines.

Contributor @[stackoverflow](#), @[jovian](#) and @[PyTorch](#) communities.

WORK EXPERIENCE

Senior Data Scientist - [British Telecom](#) (BT)

(May 2022 - Present)

- Architected and led delivery of an enterprise-grade conversational AI system using LLMs and RAG, enabling natural language queries at scale and driving a 70% reduction in manual data extraction time, while mentoring a cross-functional team across design, deployment, and production rollout.
 - Built and scaled a multimodal data pipeline using **AWS Textract**, **pdfplumber**, and **OpenSearch** to process **100K+** documents (PDFs, images, scanned files) with 90%+ accuracy, integrating AWS Bedrock for seamless LLM interactions and real-time inference with **p50/p95/p99 latency** tracking.
 - Designed and implemented a multistep agentic workflow using **CrewAI** and **LangGraph**, integrating tool-augmented pipelines for batch processing with **JSON schema validation**, retry loops on malformed LLM output, degradation – applying guardrails to minimize hallucinations in production.
 - Developed an **LLM evaluation framework** leveraging the [Ragas framework](#) and LLM-as-judge pipelines to assess faithfulness, toxicity, bias, and hallucination detection for qualitative scoring.
 - Deployed solutions on AWS with containerisation (Docker), integrated CI/CD pipelines via GitLab, and implemented tracking with **Langfuse** observability for end-to-end LLM tracing and token usage.
- Designed an automated email intelligence pipeline processing 6,000+ weekly escalation emails, eliminating manual effort and enabling real-time ticket creation in ServiceNow.
 - Built end-to-end automation using **AWS Lambda**, **Step Functions**, and **Microsoft Graph API**, leveraging **AWS Bedrock** for extracting BT product details and other relevant fields.
 - Fine-tuned **LLaMA-2 7B** locally using **QLoRA** (r=16, 4-bit NF4 quantization) on 3,000+ annotated BT escalation emails, achieving 40% improvement in field extraction F1 score over baseline prompting – deploying custom LoRA adapters to AWS Bedrock for production inference.
- Engineered a suite of recommendation systems to drive revenue growth across BT's product portfolio. Built a multi-label model using **Random Forest** and **XGBoost** to identify and recommend premium products (SD-WAN), increasing sales by 10%. Implemented a **market basket analysis** pipeline using

Apriori algorithms to uncover customer purchasing patterns for **Value-Added Services (VAS)**, achieving a 30% uplift in VAS sales – directly strengthening BT's market presence and customer retention.

Data Scientist - [Royal Dutch Shell](#)

(Sep 2016 - May 2022)

- Developed and evaluated predictive maintenance models (XGBoost, Random Forest) using SHAP-based interpretability and ROC-AUC scoring for oil refinery equipment, achieving 30% reduction in maintenance costs and 25% reduction in unplanned downtime – significantly improving workforce efficiency.
- Engineered end-to-end data pipelines in Python (Pandas, NumPy, Matplotlib, Seaborn) and developed Power BI dashboards to forecast materials on-time delivery across 5 geographies, delivering 10% budget savings within the first year of deployment.
- 5+ years of Experience with databases and big data analytics technologies in the industry including **Data Warehousing/ETL** and **relational Databases**.

IT Analyst - [TCS](#)

(Dec 2010 - Aug 2016)

- Performed **System Integration Testing & User Acceptance Testing (UAT)** to validate client PoS system.
- Spent one year in the UK guiding offshore teams with the implementation of new PoS software.
- Experience working with card and payment systems, PCI standards and ISO 8583 protocols.

Research Thesis ([View](#))

Investigated the integration of contextual language models (**GPT-3.5, Mixtral, Llama 3.1**) with classical ML models (XGBoost, Random Forest) to enhance predictive accuracy. Developed a methodology to convert tabular data into enriched textual representations and applied PCA for input optimization. Highlighting potential of LLMs in addressing data imbalance.

OPEN SOURCE DATA SCIENCE PROJECTS

AgenticFS – Autonomous Local File Organizer and Multi-Agent Orchestrator([github](#))

- **Architected and deployed agentic AI architectures for autonomous workflows**, leveraging advanced agentic frameworks, tool-using systems (CrewAI), and local LLMs (Ollama).
- Designed a **multi-agent pipeline** to partition folder classification tasks across, maximizing local LLM inference throughput and scalability via concurrency with threadpool executor.
- Engineered a **Human-in-the-Loop** safety gateway with transaction logs to prevent data loss, integrating a CLI verification step and dry-run preview mode before executing file movements.

Antigravity Studio: Autonomous Self-Correcting UI Designer([Live app](#))

- **Architected an autonomous, self-correcting agentic workflow** leveraging multimodal foundation models (**Gemini 2.5 and Imagen 4 Ultra**) to translate high-level design to visual mockups.
- **Designed a real-time streaming pipeline** to trace the AI agent's step-by-step thinking, self-correction process, and automated metrics (brand consistency, color alignment, and layout accuracy)

EDUCATION

[LIVERPOOL JOHN MOORES UNIVERSITY](#)

2023 - 2025

M.S. in Machine Learning & Artificial Intelligence

[International Institute of Information Technology Bangalore](#)

2022 - 2023

Executive Post Graduation in Data Science and Artificial Intelligence

Courses: Statistics & Probability, Python, Machine Learning, NLP, Neural Networks, MLOps

[New Horizon College Of Engineering, Bangalore, India](#)

2006 - 2010

Bachelor of Engineering, **Electronics and Communication** from **Visvesvaraya Technological University**.
