Ashwin N Hebbar

+91 - 86607 22846 | Bengaluru, IN | www.aswn.me | ashwinhebbar2k02@gmail.com | github.com/AshwinHebbar314 | linkedin.com/in/ashwin-hebbar

Software Engineer and Data Scientist with end-to-end experience building and deploying full-stack AI/ML solutions to solve complex business problems. Proven ability to take complete ownership of the project life cycle, from stakeholder engagement and data engineering to model development (LLMs, GenAI, CV) and maintaining data pipelines in cloud environments (AWS, GCP). Excels at communicating actionable insights to technical and business teams.

Skills

- **Languages & Databases:** Python, SQL, R, JavaScript, Java, C, PostgreSQL, Redis, SQLite
- **Frameworks & Tools:** FastAPI, Flask, Next.js, Apache Kafka, Git, Superset, Tableau
- **Cloud & MLOps:** AWS (Bedrock, EC2, Lambda, S3), GCP (Vertex AI, Dataproc, GCS), Azure, MLflow, DVC, Docker, Kubernetes, CI/CD
- **AI & Data Science:** LLM Fine-Tuning (PEFT, QLoRA), GenAI, NLP, Computer Vision (YOLOv8), PyTorch, Pandas, NumPy, Statistical Modeling, Feature Engineering, Big Data Pre-processing, HuggingFace

Work Experience

Data Scientist — **Dec 2025 — Present**
London Stock Exchange Group (LSEG) — *Bengaluru, IN*

- Deployed an LLM-based multilingual transliteration pipeline for 24,000+ named entities, **saving an estimated 284 person-days** of manual effort within 1.5 weeks of deployment.
- Built and owned a unified Market Processor API handling document conversion and concatenation across 6+ global markets and 5+ formats (Excel, Word, XHTML), achieving 99%+ success rate across **14,000+ files and 5,200+ document sets.**
- Engineered end-to-end financial document automation pipelines for 14+ stock exchanges globally, including sourcing agents, a parallelized Excel to PDF converter (40x speedup), and delivering ~10,000 documents to product in a single quarter.

Product Engineer, AI — **Sep 2025 — Nov 2025**
Falco Peregrinus Technologies Pvt. Ltd — *Bengaluru, IN*

- Spearheaded the development of CaseCrumbs™, a flagship interactive Business Case Study Simulation, Development and Authoring Platform. Developed AI Agents to generate 550 high quality Case Studies.
- Developed SustainSkills™, a parameterized content authoring platform, reducing workflow time from 5 hrs to 15 mins and enabling SMEs to produce 500 questions in one week, increasing **operational efficiency**.

Agentic AI Intern - Learning Innovations Team — **Mar 2025 — July 2025**
USDC Global — *Bengaluru, IN*

- Architected a GenAI data pipeline using AWS Bedrock to process and translate over 300 unstructured course PDFs, **achieving >97% semantic accuracy** and preserving complex content like LaTeX formulas.
- Pitched and demoed prototypes to senior leadership, including the VP of Learning Innovation, effectively communicating technical solutions and project outcomes to secure company-wide adoption.

Project Trainee - Space Astronomy Group — **Aug 2024 — Jan 2025**
Indian Space Research Organisation (ISRO) — *Bengaluru, IN*

- Engineered a Data Platform to process and serve **14 TB** of AstroSat satellite archive to the global research community. Earning an ‘outstanding’ recognition from ISRO Mission Scientists.
- Architected a data pipeline (**Python, Pandas, Astropy**) that **slashed complex query times by over 95%.**
- Implemented a custom caching strategy, **cutting repeat query latency by 99%** (15+ mins to 2 seconds).

Education

Program	Institution/Board	%/CGPA	Year
B. S. (Data Sci. and App.)	Indian Institute of Technology Madras, Chennai	**7.7/10**	2021 — 2025
M.Sc. (Data Science)	Chanakya University, Bengaluru	**9.5/10**	2023 — 2025
B. Sc. (CS, Math & Stat.)	Kristu Jayanti College, Bengaluru	**8.02/10**	2020 — 2023

Note: I pursued a double degree (i.e. B.S. and M.Sc.) Simultaneously.

Projects

Automating Subjective Grading with AI (FlexAGE) (link) **April 2025**

Python, GenAI, Next.js, FastAPI, PostgreSQL

- Architected and developed a full-stack, AI-powered assessment prototype to automate the grading of subjective work, **reducing evaluation time from days to seconds** for large student cohorts.
- Engineered a **rubric-grounded prompting system** using Google's Generative AI to provide fair, consistent, and context-aware feedback, solving a key challenge in scalable subjective assessment.
- Designed and implemented a secure, asynchronous backend using FastAPI and PostgreSQL, featuring a three-tiered Role-Based Access Control (RBAC) system.

Multilingual Sentiment Analysis using LLMs (link) **Feb 2025**

Python, PyTorch, LLMs, Hugging Face, W&B

- Achieved **Top 5% finish (11/223)** in a Kaggle competition by fine-tuning a Llama 3.1 8B model, boosting F1-score by 77% over baseline for 11 Indian languages.
- Applied MLOps best practices by orchestrating 29+ experiments with Weights & Biases and leveraging PEFT/QLoRA with Unsloth to **reduce VRAM usage by 50%**, enabling training on a single GPU.

Real-time Airspace Visualizer & Analytics (link) **Feb 2025**

Python, Apache Kafka, PostgreSQL, Superset

- Engineered a real-time data pipeline using Apache Kafka to stream live flight data, performing ETL in Python and loading into PostgreSQL for analysis.
- Developed an interactive Apache Superset dashboard to visualize real-time airspace traffic and key operational metrics, demonstrating proficiency in data visualization and dashboarding.

Taxi Fare Prediction using Machine Learning (link) **Nov 2023**

Python, Scikit-learn, Pandas, NumPy

- Achieved **Top 9% rank (63/714)** in a Kaggle competition by successfully predicting taxi fares with an R^2 value of **94.6%**.
- Executed a complete data science workflow, including exploratory data analysis (EDA), preprocessing, imputation for missing values, and robust feature engineering.
- Developed a reusable framework in Python to systematically test and evaluate multiple ML estimators (e.g., Linear Regression, Random Forests, Extra Trees) for optimal model selection.

Leadership, Publications & Awards

- **Best Project Award** (Jan 2024): For developing a full-stack eLibrary application with caching, async jobs, mailing, and RBAC.
- **Best Capstone Project Award** (Sep 2022): For a project improving manufacturing quality control via statistical analysis for Sealtite Dichtungs Pvt. Ltd.
- **Teaching Assistant:** Mentored a cohort of 30+ students in Linux shell scripting and data wrangling with AWK and sed at IIT Madras.
- *A Hand-Gesture Recognition System Using Image Processing to Translate Indian Sign Language Alphabets to Text (2023)* (link) and *Electric Field Mapping of Any Given Point Charge Distribution using Python (2023)* (link) in Indian Journal of Natural Sciences (Feb 2023 Issue)
- *Leveraging Machine Learning Models for Diabetes Prediction* (link), presented at "International Conference in Computational Intelligence - Feb 2025"