

DAKSH SHARMA

Delhi, India | +91 9870237700 | sharmadaksh2001@gmail.com | linkedin.com/in/sharmadaksh

PROFESSIONAL SUMMARY

AI-native Data Engineer with 2+ years of experience building cloud-native ETL/ELT pipelines, orchestration workflows, and self-serve analytics solutions across FMCG and retail operations. Hands-on with Snowflake, SQL, Python (pandas, numpy), Streamlit, and Git, with foundational knowledge of dbt and Airflow. Proven track record of delivering ingestion, transformation, and validation pipelines on modern data platforms including Data Lake and Lakehouse architectures.

WORK EXPERIENCE

Data Analyst (Intern) | DS Group, Noida, India

Apr 2026 – Present

- Built and maintained cloud-native Snowflake ETL/ELT pipelines using SQL and Python for ingestion, transformation, and validation of structured and semi-structured operational and S&OP datasets, ensuring data quality and source-of-truth reporting outputs.
- Designed and implemented event-driven continuous loading pipelines using Snowpipe, Streams for Change Data Capture (CDC), and automated Tasks with Merge statements; built JSON parsing views and dimensional for analytics-ready datasets.
- Developed AI-powered Streamlit analytics applications integrated with Snowflake datasets and Cortex; designed structured prompts (YAML/Markdown) to encode repeatable workflow logic and contributed to technical documentation and quality assurance for data models and AI workflows.

Business Support Associate | Daybreak (Musgrave Group), Dublin, Ireland

Feb 2025 – Feb 2026

- Sole BI and data analytics expert across 8 retail stores; engineered automated ETL pipelines using SQL and Power Query integrated with Power BI, eliminating 8+ manual reporting hours per week and improving supplier order accuracy by 20%.
- Designed and delivered self-serve KPI dashboards for MD and store managers, translating business requirements into scalable data models and reusable reporting workflows with strong attention to data quality and accuracy.
- Drove end-to-end analytics delivery: gathered requirements, authored functional specifications, facilitated UAT, and managed change, partnering directly with finance and operational stakeholders across the project lifecycle.
- Conducted root cause analysis and built automated data cleaning and validation scripts; tracked all delivery in JIRA under Agile/Scrum, maintaining SLA commitments and improving stakeholder trust in data outputs.
- Identified process gaps in ordering, inventory, and supplier workflows; modelled and recommended improvements that reduced overstocking by 15%, demonstrating a bias for action and outcome-driven delivery.

Operations Supervisor | ACCHL Limited, Dublin, Ireland

Jan 2024 – Jan 2025

- Built BI dashboards and ad hoc reporting using Python and optimised SQL views for a B2B/B2C wholesale business with €500K+ annual transactions; performed EDA and statistical analysis on 100K+ POS transactions to build predictive analytics and forecasting models, contributing to an 18% reduction in stock-outs and 12% reduction in operational waste.
- Maintained semantic data foundations and reusable SQL views for ongoing self-serve analysis; documented workflows, use cases, and process gaps across ordering, receiving, and inventory control.

KEY PROJECTS

SnowOps — Snowflake Health & Cost Optimisation Tool

- Built a production Streamlit application on Snowflake's ORGANISATION_USAGE and ACCOUNT_USAGE schema featuring a multi-tab dashboard covering warehouse credit consumption, query health, and storage activity; flagged 1,000+ failed and slow queries with per-user drill-down and daily automated email alerts via a stored procedure.
- Engineered a recommendations engine with HIGH/MEDIUM/LOW severity flagging to surface cost leakage points across compute, storage, and query workloads, providing admins actionable optimisation decisions with current state and recommended actions.

AI Workflow & Prompt Engineering — Internal Analytics Tool

- Designed structured prompts in YAML and Markdown to encode repeatable analytics workflows, producing reliable outputs that handle edge cases and domain ambiguity for non-technical users.
- Built reusable agentic workflow patterns integrated with Snowflake datasets and Streamlit front-ends; validated AI-generated outputs before stakeholder use.

NLP & Sentiment Analysis — MSc Thesis (Social Media Data)

- Performed large-scale EDA on 10,000+ tweets and built an ensemble ML model (BERT, RoBERTa, VADER) achieving 84% classification accuracy, outperforming all baseline models.

TECHNICAL SKILLS

- Languages:** SQL (advanced), Python (pandas, numpy, scikit-learn), R
- AI & Prompt Engineering:** Snowflake Cortex, Structured Prompt Design (YAML, Markdown), Agentic Workflow Design, NLP, Transformer Models (BERT, RoBERTa), PyTorch, TensorFlow, XGBoost, Orchestration Frameworks
- Data Platforms & Cloud:** Snowflake, PostgreSQL, MongoDB, BigQuery, AWS (S3), GCP, Oracle
- Data Engineering:** ETL/ELT Pipelines, Data Modelling, Dimensional Modelling, Data Validation, dbt (foundational), Airflow (foundational), PySpark, Databricks, Design patterns, Data processing
- BI & Visualisation:** Streamlit, Power BI, Tableau, Advanced Excel (Power Query, Pivot Tables, VBA, Macros)
- Analytics:** Self-Serve Analytics, Semantic Data Modelling, EDA, Statistical Analysis, Forecasting, Predictive Analytics, A/B Testing, KPI Reporting, Root Cause Analysis, Financial Planning, Problem solving
- Tools & Methodologies:** Git, JIRA, Agile/Scrum, CI/CD, User testing, Requirements Gathering, Functional Specifications, Process Mapping, Gap Analysis, Stakeholder Management, Workflow Documentation, ERP, QA

EDUCATION

Master of Science — Data Analytics

2023 – 2024

National College of Ireland, Dublin | First-Class Hons. (1.1)

Bachelor of Science — Statistics (Honours)

2019 – 2022

University of Delhi | GPA: 9.243 / 10

CERTIFICATIONS & PROFESSIONAL DEVELOPMENT

Snowflake Data Cloud Certifications (Badges 1–6, Accelerate Series, Northstar Builder Series) | SAP Datasphere & Generative AI Complete Python Bootcamp | Databricks Lakehouse Fundamentals