

UDISHA KAURA

☎ 438-778-0693 ✉ udishakaura@gmail.com [in](#) [LinkedIn](#) [G](#) [github](#) [Portfolio](#)

Data Engineer with 3+ years of experience designing and maintaining Azure data lakehouse solutions across enterprise environments. Proficient in the full Microsoft Azure stack Databricks, Data Factory, Synapse, Fabric, and Purview with hands-on expertise in PySpark, Python, and advanced SQL for high-volume ETL/ELT pipelines.

Education

Master of Engineering in Software Engineering

Concordia University

Sep. 2023 – April 2025

Montreal, Quebec

Bachelors of Technology in Computer Science

Guru Nanak Dev University

June. 2018 – June 2022

India

Technical Skills

Data Engineering & Cloud: Azure Data Factory, Azure Databricks, Azure Synapse, Microsoft Fabric, Delta Lake

Programming & Querying: SQL (Expert), Python (Advanced), PySpark, Pandas

BI & Analytics: Power BI (Datasets, DAX, Data Modeling, RLS), Ad-hoc SQL Analysis

Data Modeling & Quality: ETL/ELT Design, Data Quality Testing & Validation, Data Lineage, Metadata Management, Lakehouse Architecture

Governance & Collaboration: Azure Purview, Data Governance, Azure DevOps, CI/CD Pipelines, GitHub Actions

Experience

Data Engineer (Techo-Bloc)

March 2025 – Present

Databricks, PySpark, SQL, Python, Azure Data Factory, Power BI,

Montreal

- Architected and deployed 20+ end-to-end Medallion Architecture pipelines using **Databricks (PySpark/SQL)** and Delta Lake, improving enterprise reporting speed by 40% through DirectQuery optimization.
- Orchestrated complex data workflows using **Azure Data Factory (ADF)** and **Synapse**, maintaining 99.9% pipeline uptime and ensuring 95%+ on-time availability for mission-critical executive dashboards.
- Led the data migration strategy from Legacy AX to Dynamics 365 (D365), re-engineering relational logic for 12+ core KPI datasets and ensuring 100% data parity during the transition, while establishing data stewardship standards for Gold-layer datasets.
- Optimized Spark query performance through partitioning, Z-Ordering, and broadcast joins, reducing compute costs by 25% and cutting troubleshooting effort by 30 hours per month.
- Registered Delta Lake assets in **Azure Purview** for metadata cataloging and lineage tracking across Medallion layers.
- Evaluated and piloted **Microsoft Fabric** notebooks and Fabric Data Factory pipelines as part of a planned migration from ADF, documenting architecture tradeoffs for stakeholder review.

Data Engineer Intern (Cozey)

May 2024 – August 2024

Python, SQL, Zendesk API, Pandas, ETL Pipelines, Data Validation, Git

Montreal, Quebec

- Developed automated ETL pipelines in Python to integrate Zendesk API data, including tickets, chats, and transcripts, centralizing customer sentiment data for the product team.
- Responded to 15+ **ad-hoc SQL** requests from the product and operations teams, querying Zendesk and transactional data to surface trends in ticket volume, resolution patterns, and customer friction points.
- Built a Customer Service KPI suite covering resolution time, and volume trends improving visibility into support bottlenecks and contributing to a 12% improvement in response times.
- Improved data reliability by implementing automated validation scripts for schema drift detection and duplicate handling in Git-versioned workflows, reducing downstream reporting errors by 20%.

Analytics Engineer (KPMG)

January 2022 – August 2023

SQL, ETL/ELT, Stored Procedures/Views, Data Governance.

Gurugram, India

- Standardized ELT transformation logic across 8+ client engagements by authoring reusable stored procedures and SQL templates, cutting onboarding time for new data migration projects by 45% and reducing ad-hoc scripting.
- Translated complex business reporting requirements from cross-functional stakeholders into SQL views and transformation workflows, processing 15M+ records weekly and delivering investment-grade data to BI consumers with consistent SLA adherence
- Diagnosed and resolved query bottlenecks in high-concurrency reporting environments by restructuring joins, rebuilding indexes, and refactoring layered views — cutting average query latency by 50% across multiple client dashboards.

Certifications

- Databricks Data Engineer Associate Preparation
- Google Business Intelligence Professional Certificate
- AZ-900 Azure Fundamentals Preparation
- Google Data Models and Pipelines Certification