

Shrey Niraula

+1 (475) 251-5056 nshrey53@gmail.com [Linkedin](#) [Github](#) [Portfolio](#)

Computer Science graduate with 4+ years of experience specialized in AWS architectures and Agentic AI frameworks. Proven ability to bridge the gap between Big Data infrastructure and LLM-driven automation, including developing autonomous cloud optimization agents and high-performance ETL pipelines.

PROFESSIONAL EXPERIENCE

NASA IMPACT (CSDA Program) | *Data Engineer*

Aug 2024 – Present

- Architected an Agentic Cloud Decision Framework utilizing self-hosted LLMs (Ollama) to perform tool-augmented reasoning over petabyte-scale storage inventories, engineered a code-interpreter agent to automate multi-tier storage recommendations between MCP and NGAP while incorporating Human-in-the-Loop (HITL) checkpoints for final governance.
- Migrated legacy vendor ingestion workflows (Airbus and BlackSky) into a partitioned Airflow pipeline using PySTAC for metadata standardization, ensuring end-to-end reliability through integrated Pytest validation suites for expansion logic and transformations.
- Automated the manual MAXAR checksum verification workflow with an Airflow-orchestrated solution that exports metadata from DynamoDB and detects missing checksums via Athena, eliminating 100% of manual validation effort.
- Migrated legacy Airflow DAGs to modern standards and optimized key workflows, reducing pipeline runtime.

Tech Stack: Apache Airflow, AWS (Athena, DynamoDB, S3, Glue), PyStac, Metadata mapping, Pytest

GrowByData Services | *Mid-Level Data Engineer*

May 2022 – July 2024

- Owned end-to-end integration from POC, acquisition, ETL, to reporting backend logic for LeadGen and Amazon data sources.
- Re-architected and scaled Pinterest, Facebook, TikTok, SEO, Google Ads, Bing Ads, and Shopify ingestion pipelines using multi-threading and memory optimization, achieving 2–5× faster acquisition speeds.
- Migrated monolithic Redshift architecture into a lakehouse-based system (S3 + Glue Catalog + Dremio), reducing compute cost by over 2× and eliminating CPU bottlenecks.
- Refactored existing data anonymization logic using Python decorators, reducing demo report development time (100+ reports) from 3 months to 2 weeks.
- Integrated Redis caching into the reporting layer, reducing query latency by around 70%.
- Mentored junior engineers and trained interns in Python, SQL, Spark ETL, dimensional modeling, lakehouse architecture & so on.

Tech Stack: Lakehouse, Datalake, EMR, Glue, Crawler, Redis, Optimization, ETL Orchestration, Dremio

GrowByData Services | *Associate Data Engineer*

May 2021 – May 2022

- Built the first production acquisition pipeline for Pinterest Ads with OAuth2 authentication.
- Scaled existing Facebook Ads acquisition to handle throttle-limit issue, ensuring stable acquisition.
- Resolved production issue in ETL pipelines, gaining hands-on experience with star-schema modeling and Slowly Changing Dimension (SCD).
- Migrated Talend ETLs into distributed Spark (Scala + SparkSQL), cutting Redshift CPU load by 2-3x and reducing ETL latency by around 3x times.

Tech Stack: Data Warehousing, Data Modeling, Talend, PySpark, Scala, SparkSQL, Redshift, Star-Schema Design, OAuth2

SKILLS

- **Programming Languages:** Python, Scala, SQL
- **Cloud & Big Data Platforms:** AWS (EC2, S3, Athena, Glue, DynamoDB, EMR, Lambda, MSK), Spark, Apache Airflow
- **Databases & Caching:** MySQL, Postgres, Dremio, Iceberg, Redis
- **Software Engineering:** Flask, Docker, RabbitMQ, Git, CI/CD, OOP, RESTful APIs, Multithreading, Terraform, Distributed Systems
- **Data Concepts:** Data Lake, Data Warehousing, Data Modeling, SCD, Data Orchestration, Data Governance, ETL/ELT, OLAP
- **AI/ML:** Local LLMs, Agentic Frameworks (LangGraph), RAG (Retrieval-Augmented Generation), Prompt Engineering, Vector Databases.
- **Data Science:** Statistical Modeling, Pandas/NumPy, Matplotlib

EDUCATION

The University of Alabama in Huntsville

Huntsville, AL, USA

Master's in Computer Science - Data Science Concentration - GPA 4.0

Aug 2024 – Present

- Key Subjects: Advanced Database, Data Structure and Algorithm, Big Data Computing, Cloud Computing, Machine Learning
- Thesis: Agentic Cloud Decision Framework: Multi-agent local LLM orchestration and tool-augmented reasoning to provide recommendation on large-scale NASA cloud storage.

Tribhuvan University, Institute of Engineering, Pulchowk Campus

Lalitpur, Nepal

Bachelor's in Electronics and Communication Engineering - 81.70% (Rank: 4th Position, top 1%)

2016 – 2021

- Ncell Excellence Cash Award (2019) for being consecutive 3rd & 4th semesters topper
- Key Subjects: Object Oriented Programming, Database, Artificial Intelligence, Data Mining, Linear Algebra, Computer Architecture

PROJECTS

Real Time Visual Localization and Mapping of Mobile Robot in Dynamic Environment | *Final Year Project*

2020

- Visual SLAM based project to localize and navigate the mobile robot in dynamic environment.
- Fine-tuned semantic segmentation model (ICNet) to detect the human, considered as dynamic obstacle so that robot does not consider them as valid landmark while localizing itself.
- **TechStacks:** Semantic Segmentation, PyTorch, Model Fine Tuning

Vehicle Traffic Management and Analysis | *Semester Project*

2020

- Project to synchronize the traffic lights across three junctions of Kathmandu
- Developed web app to simulate the traffic junctions. The app could communicate and control junctions like a control center.
- **TechStacks:** Python Flask web app, JS, AJAX/JQuery, Data Collection and Simulation, Data Modelling

CERTIFICATIONS

- **EDX (2020):** *Applied Deep Learning Capstone Project (DL0320EN)* [↗](#)
- **Coursera (2020):** *Machine Learning* — Stanford University (Andrew Ng) [↗](#)
- **Coursera (2020):** *Deep Learning Specialization* [↗](#)