

Mayur Dalvi

mayurdalvi.5@gmail.com | 720-335-3625 | [linkedin.com/in/mayur-dalvi](https://www.linkedin.com/in/mayur-dalvi) | github.com/mayurdalvi | Willing to Relocate

TECHNICAL SKILLS

Core Skills: Python, R, SQL, SAS, Excel (Pivot Tables, VBA), Salesforce CRM, Power Platform, PL/SQL, SSIS, JIRA, Linux, MS-SQLServer, Snowflake, PostgreSQL, Oracle, DBT, GitHub, Hadoop, Databricks

ML & Cloud Tools: Pandas, NumPy, Scikit-learn, Regression, Time Series Analysis, A/B Testing, CI/CD, MLOps, PySpark, PyTorch, AWS (S3, Lambda, SageMaker, Redshift), Azure ML Studio, GCP BigQuery

BI Tool: Power BI (DAX), Tableau, Alteryx, Qlik, Looker, Matplotlib, Seaborn, EDA

Data Management: Airflow, Data Warehouse, Data Modeling, Data Governance, Data Integration, Data Migration

Skills: Data Analysis, Forecasting, Marketing Analytics, Risk Management, Project Management, Agile, Data Visualization

PROFESSIONAL EXPERIENCE

City and County of Denver | Data Analyst | Denver, Colorado, USA

December 2024 - Present

- Automated audit workflows across 217 city locations using Python and SQL, resulting in a 65% reduction in manual effort and improved fraud detection accuracy by implementing anomaly detection models and data analytics dashboards.
- Streamlined procurement monitoring by integrating Workday ERP data models with geospatial analysis and validation logic, ensuring 100% policy adherence across 122K+ purchase orders and enhancing transparency.
- Engineered and deployed forecasting models in Python using Prophet and XGBoost, improving revenue prediction accuracy by 28%, enabling data-driven decision and optimizing budget allocations across 15 city departments.
- Built and fine-tuned a Snowflake Document AI (Arctic-TILT) model to extract key invoice data from 10K+ documents with 95% accuracy, automating data ingestion, processing, and storage via streams, tasks, and SQL procedures.

Deque Systems | Data Analyst | Boulder, Colorado, USA

January 2025 - May 2025

- Collaborated with cross-functional teams to analyze engagement of 10K+ Deque University learners using cohort-based retention models in SQL and Python, identifying modules with 18% higher completion rates to guide curriculum updates.
- Developed interactive dashboards for leadership and cross-functional stakeholders, tracking course adoption, and learner progression insights contributed to a 25% increase in course completion and 15% improvement in subscription renewals.

ICR Inc | Data Scientist | Boulder, Colorado, USA

August 2024 - December 2024

- Automated ingestion and preprocessing 10 years of aviation and weather data using RESTful APIs, and BeautifulSoup, reducing data preparation time by 60% while maintaining data integrity and scalability for predictive analytics modeling.
- Improved flight delay prediction accuracy by 20% by retraining XGBoost models (ROC-AUC: 0.80), conducting hyperparameter tuning in Scikit-learn, and delivering actionable insights to optimize airline operations and scheduling.

NICE | Data Engineer | Pune, India

July 2021 - July 2023

- Developed and optimized scalable ETL pipelines in Python and SQL to ingest and validate 10TB+ of call data, implementing data integrity checks that improved data quality and reduced system downtime by 25%.
- Automated batch extraction and transformation workflows with REST APIs and MySQL database, increasing data throughput and accuracy by 20%, while eliminating manual dependencies in daily operations.
- Improved pipeline execution speed from 12 to 8 hours (33% faster) by implementing parallel processing, advanced exception handling, and system resource tuning, ensuring scalable, high-performance data infrastructure for enterprise workloads.
- Designed and deployed end-to-end data pipelines on AWS using Spark and Redshift, enabling real-time advanced analytics on 5B+ records, reducing reporting latency by 60% and supporting scalable, cost-efficient business intelligence operations.
- Consulted NICE Cloud customers, enhancing data retrieval through NICE Data Explorer (NDE) via RESTful APIs, achieving a 40% reduction in data retrieval time, enhancing overall system efficiency.

PROJECTS

Fraud Detection | Python, PowerBI, Excel | [Link](#)

- Built an end-to-end fraud detection system using Python, XGBoost, and Power BI, improving model precision by 18% and reducing investigation time by 25% through feature engineering, hyperparameter tuning, and real-time KPI dashboards.

EDUCATION

Master of Science in Data Science

August 2023 - August 2025

University of Colorado Boulder, Colorado, USA

GPA: 3.86/4.0

Relevant Courses: Machine Learning, Data Mining, Natural Language Processing, Mathematics, Statistics, Business intelligence