

Alex Morgan

Austin, Texas | alex.morgan.dev@example.com | (555) 010-1234 | linkedin.com/in/alex-morgan-data | github.com/alexmorgan-dev

EXPERIENCE

Northwind Analytics

Austin, USA

Data Engineer

Jan 2024 – Present

- Built PySpark/AWS Glue pipeline ingesting 3.5B+ daily records into Delta Lake, powering analytics, ML features, and reporting for 12 downstream teams – structured on Medallion Architecture
- Developed internal Lambda API tooling that automated pipeline configuration, testing, and deployment workflows, reducing engineering effort by ~70%
- Engineered serverless ingestion monitoring via Lambda, DynamoDB, and Step Functions (8-state workflow), automating delta file event triggering and improving pipeline observability
- Automated end-to-end test data provisioning by chaining 10+ Python and shell script steps across internal APIs and DynamoDB, cutting setup time ~75% and eliminating manual intervention
- Diagnosed and resolved 15+ high-severity data pipeline incidents via root cause analysis, eliminating recurring SLA breaches across 12 downstream data consumers
- Delivered 10+ Delta Lake table integrations within strict delivery SLAs by aligning requirements with cross-functional downstream teams and securing stakeholder sign-off, authoring SQL transformation configs and schema validation logic in Databricks

BlueOrbit Labs

Remote, USA

Data Engineering Intern

May 2023 – Aug 2023

- Designed an Airflow + dbt ETL pipeline loading 50M+ daily events from Kafka into Snowflake, cutting data-availability latency from 6 hours to under 30 minutes for 5 analytics teams

Acme AI

Remote, USA

Machine Learning Engineer Intern

May 2022 – Aug 2022

- Led team of 3 to build NLP preprocessing pipelines (TF-IDF vectorization, GloVe tokenization, BERT tokenization) for 40,000+ imbalanced support tickets across 5 intent classes
- Benchmarked Naive Bayes (TF-IDF), LSTM (GloVe 300-dim), and BERT (110M params) on 5-class intent classification, achieving 95.4% validation accuracy using TensorFlow and scikit-learn

SKILLS

Programming & Tools: Python, SQL, Linux/Unix, Git, Jenkins, Docker, Terraform, GitHub Actions

Databases: MySQL, Snowflake, MongoDB, PostgreSQL, Neo4j, BigQuery

Big Data & Visualization: PySpark, Spark, Delta Lake, dbt, Databricks, Airflow, Kafka, Hadoop, QuickSight

ML & AI: PyTorch, scikit-learn, HuggingFace Transformers, MLflow, BERT, LSTM, CNNs

AWS Services: Lambda, Glue, S3, DynamoDB, Redshift, Kinesis, IAM, EMR, Step Functions, Athena, EC2

PROJECTS

MarketPulse: Real-Time Streaming Analytics Platform

Jan 2024 – Mar 2024

- Built a Kafka + Spark Structured Streaming platform on AWS (MSK, EMR, S3, Redshift) processing 200K+ events/sec, surfacing sub-second operational dashboards in QuickSight with automated data quality checks on every micro-batch

PUBLICATIONS

Adaptive Pooling Functions for Convolutional Neural Networks | *IEEE Example Conference 2022 – Best Paper Award*

- Implemented a novel adaptive pooling aggregation in CNN layers using PyTorch; outperformed classical pooling on state-of-the-art architectures across benchmark image datasets

AWARDS & HONORS

Riverside State Datathon 2023: Secured 1st place, winning \$500 for predictive demand modeling of retail sales

Graduate Merit Fellowship: Awarded \$15,000 tuition fellowship in recognition of academic excellence

EDUCATION

Riverside State University

Austin, TX, US

Master of Science in Data Science | GPA: 3.9/4.0

Aug 2022 – May 2024

Westbridge Institute of Technology

Remote, US

Bachelor of Technology in Computer Engineering

Jul 2018 – May 2022