

Landon Townsend

Clarksburg, MD | (301)-580-4469 | [LinkedIn](#) | [GitHub](#) | landontownsend20@gmail.com

EDUCATION

University of Maryland - College Park

Master's of Science, Applied Political Analytics (Data Science & Automation Focus)

Relevant Coursework: Advanced Coding & Automation, Research Design, Advanced Data Display

Bachelor's of Science, Government & Politics

Cumulative GPA: 3.73

College Park, MD

Expected May 2027

May 2026

WORK EXPERIENCE

Data Engineer Intern

Baltimore Aircoil Company

Jessup, MD

June 2026 - August 2026

- Architected and developed a full-stack Power BI embedded analytics portal (Python/FastAPI, Azure AD, Microsoft Fabric) consolidating 3,000+ reports across business workstreams, delivering centralized analytics access to a 5,000 person organization during an active SAP-to-Fabric data warehousing migration.
- Designed self-service report governance system enabling any of 5,000+ authenticated employees to publish Power BI reports to a shared catalog via REST API ownership verification, eliminating single-administrator bottleneck across 100+ workspaces.
- Led enterprise-wide Power BI workspace audit across 100+ workspaces as a data governance initiative, delivering structured findings to senior stakeholders.

Data Scientist Intern

Analytica

Bethesda, MD

February 2026 - May 2026

- Architected and deployed agentic AI governance framework (GSR) on Azure VM, integrating a FastMCP multi-provider LLM gateway (OpenAI, Anthropic, Ollama), PostgreSQL database with 8 targeted indexes, and a Dockerized Streamlit dashboard all configured for zero-touch auto-start via systemd.
- Built a Python-based evaluation pipeline ingesting 100% of LLM responses through 4 custom metric modules (groundedness, hallucination, adversarial detection, risk scoring), writing structured scored records to PostgreSQL with a 0-to-dashboard latency of under 30 seconds, eliminating 100% of manual quality review across all provider calls.
- Designed a relational PostgreSQL schema (4 tables, 8 indexes, 95-day retention policy) supporting concurrent multi-source data ingestion from both live MCP traffic and batch test runs, with foreign key constraints and pre-aggregated KPI metrics enabling sub-second dashboard queries at scale.

Data Engineer Intern

U.S. Department of Agriculture, Office of Analytics and Innovation (OAI)

Washington, D.C

January 2025 - May 2025

- Engineered ETL pipelines using R and SQL to process data from three disparate sources, reducing manual cleaning time by 95%.
- Architected cloud migration strategies utilizing Azure Databricks and Azure Data Factory; established automated CI/CD-style data pipelines for continuous synchronization.
- Developed reusable, modular codebases in R to standardize data manipulation, serving as the foundational infrastructure for OAI's cloud systems.

PROJECTS

Talk to Your Data - FDA Agentic RAG Architecture

Analytica

- Engineered a 3-tier RAG pipeline replacing a 22-tool keyword-matching architecture for an FDA data catalog, reducing codebase complexity by ~70% and per-query operational cost to \$0 by running ChromaDB vector search and Llama 3.2 entirely on local infrastructure.
- Embedded 39 catalog items across 5 data types (tables, glossary terms, stewards, articles, SQL queries) using sentence-transformers, enabling semantic search across 10 regulatory tables and eliminating 13 separate JSON file lookups with a single vector query.
- Delivered a full-stack POC in 2 weeks including a FastAPI backend, production-quality chat UI with agentic reasoning panel, interactive solution architecture diagram, and FDA stakeholder persona cards, built to support a leadership funding decision.

Reddit Political Discourse ETL Pipeline & Analysis

- Built ETL (Python, SQLite, Pandas) on Reddit records, applying unsupervised LDA to segment political discourse into 5 thematic clusters.
- Deployed RoBERTa (Hugging Face) for sentiment classification, leveraging transfer learning to capture advanced semantic nuance.
- Orchestrated ETL and analysis workflow as Apache Airflow DAG (Python, TaskFlow API), converting 4 manually-run scripts into 6 dependency-managed tasks with automatic retries and a monitored dependency graph, running independently.

SKILLS & CERTIFICATIONS

- **Technical:** Python, SQL, R, PowerBI, Microsoft Fabric, Databricks, Tableau, Docker, FastAPI, Apache Airflow
- **Certifications:** Microsoft Certified: Azure Data Fundamentals (DP-900), IBM Data Analyst
- **Languages:** Arabic (Professional Fluency), French (Professional Fluency)