

Mohammad Haider

Software Engineer | Data platforms, analytics systems & reliability

hi@mhaider.dev | mhaider.dev | github.com/haidmoham | linkedin.com/in/haidmoham

Software engineer with Microsoft and University of Virginia experience in data migrations, reconciliation, and analytics delivery. Builds reliable data products and reproducible evaluation workflows.

PROFESSIONAL EXPERIENCE

University of Virginia

May 2025 - Mar 2026

IT Analyst Intermediate, BI Developer

Charlottesville, VA

- Owned SQL-backed Tableau delivery from stakeholder interviews through deployment and maintenance; translated ambiguous operational questions into shared metric definitions.
- Reconciled purchase-order states across financial systems and built reporting that surfaced discrepancies and their financial impact.
- Reduced organization-wide dashboard load times by **approximately 20%** through collaborative, telemetry-driven optimization.
- Built environmental and economic-impact reporting for the ReUSE Store that reached **over 400 views**; adapted Tableau views as operational workflows changed.

Microsoft

Jul 2021 - Jul 2024

Software Engineer I

Redmond, WA

- Owned a **zero-downtime** SSAS Multidimensional-to-Tabular migration; defined parity checks, monitored **four weeks** of pipeline behavior, and kept the legacy service available until cutover criteria were met.
- Wrote and deployed Azure Data Factory verification scripts for daily pipeline monitoring and legacy-to-Tabular data consistency.
- Validated downstream reporting workflows showing **up to 200x** faster filtering and **approximately 4x** faster Excel loading.
- Built a full-stack C#/.NET prototype that guided configuration, generated configuration files, and validated output.

Microsoft

Summers 2017, 2018

Software Engineering Intern

Redmond, WA

- Built a C# service that identified each OneDrive user's most recently used items for a product UI feature; validated the service in development.

SELECTED ENGINEERING PROJECTS | INDEPENDENT

Indigo Circuit | Tournament analytics platform

indigocircuit.app

- Built and deployed a tournament analytics platform with API/scraped ingestion, DuckDB/dbt marts, and Flask dashboards; sustained updates within 24 hours of every major tournament since April.
- Reduced refresh-related read blocking from ~9 minutes to under 1 second through shadow builds and atomic publication; preserved live data through two observed production rebuild failures.

LM Lab | Language-model engineering

github.com/haidmoham/lmlab

- Completed **39 language-model training runs** across three controlled studies; persisted source/corpus hashes and model/optimizer/RNG state, verifying restoration against subsequent updates.
- Compared 8- to 256-token contexts in 12 training runs, each with **2.56 million target-token exposures**; evaluated frozen models across **30 supported train/evaluation-context combinations**.

TECHNICAL SKILLS

Languages: SQL, Python, C#, R

Production: Azure Data Factory, SSAS Tabular, Azure Synapse, Kusto, Tableau, C#/.NET

Project stack: DuckDB, dbt, Flask, pandas, NumPy, PyTorch, Jupyter, Git, GitHub Actions

EDUCATION

Virginia Tech | B.S. in Statistics

2021