

# DHWANIL MORI

Arlington, VA | [dhwanimori@gmail.com](mailto:dhwanimori@gmail.com) | [LinkedIn](#) | [Github](#) | [Portfolio](#)

## PROFESSIONAL PROFILE

---

Founder and data science leader with extensive experience designing end-to-end AI solutions for retail analytics and research platforms. Proven track record delivering a multi-agent demand-forecasting engine and building a full research infrastructure that drives actionable insights and A/B-tested improvements. Skilled in Python, SQL, TensorFlow, and scalable ETL pipelines, with strong expertise in predictive modeling and generative AI. Seeking to apply this experience to accelerate data-driven product innovation and measurable business impact.

## EDUCATION

---

### The George Washington University

Master of Science, Data Science

- **GPA:** 3.4/4

- **Achievements:** Partnered with Anthropic through RAIN, an AI-native retail intelligence platform

Aug 2024 - May 2026

Washington, DC

### Gujarat Technological University

Bachelor of Science, Computer Science (AI/ML)

- **GPA:** 9.76/10

Jan 2020 - May 2024

Ahmedabad, India

## TECHNICAL SKILLS

---

- **Languages:** Python, R, SQL
- **Machine Learning & AI:** TensorFlow, Scikit-learn, Boosted Trees (XGBoost/LightGBM), Logistic Regression, Classification, Clustering & Sampling Techniques, Unsupervised Models, Reinforcement Learning, Deep Learning, NLP, Multi-Agent LLM Systems, Generative AI (ComfyUI, MusicGen), Claude API, Propensity Modeling, Multinomial Logistic Regression
- **Statistics & Experimentation:** A/B Testing, Holdout Experiment Design, Power Calculations, Exploratory Data Analysis (EDA), Feature Engineering & Transformation, Regression Analysis, Statistical Inference
- **Data & Analytics Tools:** Tableau, Power BI, AWS QuickSight, Microsoft Excel, Jupyter Notebook, Stata, Gephi
- **Databases & Data Engineering:** MySQL, Neo4j, Hadoop, Spark, ETL Pipelines, AWS, Google Cloud, Box API
- **Infrastructure & DevOps, Developer tools:** HPC Clusters (Cerberus, Pegasus), Slurm Job Scheduling, Remote GPU Computing, Linux/Unix, SSO Authentication, n8n, VS Code, PyCharm, RStudio, GitHub, Hubspot

## PROFESSIONAL EXPERIENCE

---

### RAIN (Retail Analytics & Intelligence Network) - Backed by Anthropic

Founder

- Developed a retail operations platform using AWS and Python with a multi-agent workflow that integrates inventory automation, supplier coordination, and demand forecasting for grocers and restaurants, reducing manual work and improving stock accuracy.
- Built a multi-agent AI engine with Python, SQL, and Stata that models demand shifts, detects real-time anomalies, and provides replenishment recommendations, improving forecast accuracy and lowering out-of-stock incidents
- Conducted 30+ customer discovery interviews with independent grocery retailers and restaurant owners, identifying recurring pain points around stockouts, overstock cycles, and supplier delays that directly shaped the platform's core feature roadmap.

Dec 2025 - Present

Washington, DC

### Heartwise

Head of Research Design

- Architected research infrastructure using Jupyter Notebook and n8n, designing surveys, defining outcome metrics, and building data pipelines that combined user activity, session feedback, and completion data into dashboards and investor impact reports
- Ran A/B tests on features, session content, and matching logic with Jupyter Notebook analysis, optimizing user experience and observing growth in communication confidence and relationship outcomes
- Engineered a matching system with privacy-compliant data collection, informed consent frameworks, and transparent scoring based on lifestyle and communication preferences, using clustering techniques; the system improved match relevance and increased user satisfaction

Jan 2026 - Present

Washington, DC

### Resilience, Inc

Data Analyst

- Collected, cleaned, and validated nonprofit program data across multiple sources using Python and SQL, ensuring data integrity and readiness for downstream analysis and reporting, which reduced data-processing errors and accelerated report generation
- Conducted exploratory and trend analysis on real-world program datasets to surface actionable insights, delivering clear visualizations in Excel and Python that resonated with both technical teams and non-technical stakeholders.
- Partnered with cross-functional teams to interpret program performance data, translate findings into practical recommendations, and support ad hoc analytical projects that informed organizational decision-making.

Feb 2026 - Present

Tampa, FL

### The George Washington University, Corcoran School of the Arts and Design

Research Assistant

May 2025 - Jan 2026

Washington, DC

- Consolidated and validated data from Box API, AWS, Google Sheets, and SQL databases into structured pipelines, ensuring accurate and accessible data for non-technical users
- Enabled non-technical users, including artists and students, to run generative AI models without coding by browser-based sandbox platform with SSO authentication and streamlined access to ComfyUI and MusicGen
- Increased generative AI performance by using GW's HPC clusters for GPU processing, integrating Slurm job scheduling, and automating output delivery to Box and local storage through a pipeline that linked Google Sheets, ComfyUI, and AWS

**The George Washington University – Department of Physics**

**Sep 2025 - Dec 2025**

*Research Collaborator | Project: When AI Fails: Exposing the Hidden Dangers of Intelligent Systems*

*Washington, DC*

- Collaborated with an interdisciplinary team, employing Gephi and Jupyter Notebook to investigate failure modes, biases, and systemic risks in open-source AI models such as EleutherAI/gpt-neo-2.7B, Qwen, Gemma, and Microsoft/phi-2, producing a comprehensive risk-assessment report
- Used machine-learning diagnostics and adversarial testing on Hadoop and HPC clusters to pinpoint failure points in AI decision-making related to temperature, threshold, and other parameters, enabling the team to prioritize model refinements
- Analyzed model outputs with symbolic representation on HPC clusters, uncovering patterns in residual internal states that informed subsequent studies of AI dynamics

**George Washington University – Department of Physics**

**Nov 2025 - Present**

*Research Assistant | Project: LLM Council using multiple AI agents*

*Washington, DC*

- Designed and executed multi-agent LLM experiments in Python, modeling El Farol-style decision environments to study coordination failures across heterogeneous AI agents - findings directly informed a peer-reviewed manuscript under review at JAIAI and IEEE.
- Built pipelines to assess agent decision quality through binary labeling schemas, feedback loops, and instruction-policy adaptation mechanisms, generating reproducible insights into how LLM agents break down under competitive and cooperative conditions.
- Analyzed agent dynamics using Python, reinforcement learning algorithms, and iterative simulation frameworks, contributing findings that directly informed a peer-reviewed manuscript currently under review at JAIAI and IEEE Journal

**PUBLICATIONS**

- 
- Mori, M. Dhwanil, Prof. Neil F. Johnson. Safe and Efficient Resource Allocation in Multi-AI-Agent LLM Systems. *Under review at Journal of Artificial Intelligence and Advance Intelligence (JAIAI), 2026 | featured in Button-down AGI newsletter*
  - Mori, M. Dhwanil. Temperature-Induced Symbolic Phase Transitions in Autoregressive Language Models. *Co-author with Prof. Neil F. Johnson and 22 other members | GW CCAS Research Showcase, April 2026*
  - Mori, M. Dhwanil. Forecasting when deliveries fail before they do: a practical ML approach for procurement team. *with academic advisor Prof. Junjun Yin | GW CCAS Research Showcase, April 2026*

**AWARDS & FEATURED**

- 
- **Google Developer Startup Day Pitch Competition (2025):** Washington DC  
Won for developing Onboard AI, an intelligent automation platform for onboarding process for internal company software pitched at GWU's Build with AI Startup Day.
  - **GW's Trustworthy AI Hackathon (2026):** Washington DC  
Awarded three prizes for the Most Innovative Use of AI, the Longest AI Prompt, and the Most Pull Request. Also, our project was featured in George Washington University's Data Science Newsletter.