Ishan (Friday) Mishra

Pittsburgh, PA 15213 • imishra@andrew.cmu.edu • (608) 895-1067 • github.com/ishanmishra128 • linkedin.com/in/ishanmishra128/

EDUCATION

Carnegie Mellon University, Heinz College — Pittsburgh, PA
Master of Information Systems Management — Business Intelligence & Data Analytics (BIDA) Track — *December 2027*
- Relevant Coursework: A/B Testing & Experimentation, Intermediate Database Management, Data Engineering, NoSQL Data Management, Organizational Design & Implementation, Python for Software Development, Linux & Open Source

University of Wisconsin-Madison — Madison, WI
Bachelor of Business Administration - Information Systems & Data Science — *May 2026*
- GPA: 3.87/4.00 | Graduated with Distinction | Dean's List, 8 of 8 consecutive semesters

PROFESSIONAL EXPERIENCE

University Health Services — Madison, WI
Informatics Analyst Intern — *March 2023 - May 2026*
- Led data governance and metadata management initiatives by establishing standardized data workflow templates for 300+ providers, ensuring data quality, security, and enterprise standards compliance.
- Delivered 25+ ad hoc SQL reports monthly for non-technical clinical and administrative stakeholders, reducing reporting turnaround time by 10+ hours monthly and improving data-driven decision-making.
- Redesigned EHR-integrated medical templates and patient intake forms, raising submission completion rates to 100% and reducing data entry errors, saving clinical staff 3+ hours weekly on corrections and patient follow-ups.
- Automated master data and data operations audits using Python, executing rigorous data quality assurance processes that reduced preparation time by 85% while meeting strict HIPAA compliance standards.

Public Education ETL Pipeline | *Python, SQL, R* — Madison, WI
Data Analyst - Wisconsin School of Business Faculty-Sponsored Research — *February 2025 - May 2026*
- Engineered end-to-end ETL/ELT pipelines in Python to acquire, cleanse, and transform public education data from 500+ sources into a unified, analysis-ready dataset covering all 50 U.S. states across 6 years, improving data reliability and platform performance.
- Conducted paired t-test hypothesis testing in R to evaluate the relationship between school digital policies and student academic performance, supporting faculty-led research on learning outcomes.

TECHNICAL PROJECTS

E-Commerce Conversion & Retention Analytics Dashboard [Repo] | *BigQuery, dbt, Tableau, SQL* — *June 2026 - August 2026*
- Architected scalable cloud data pipelines and data models using BigQuery, dbt, and Tableau to process 360K+ user sessions, partnering with stakeholders to deliver fit-for-purpose data solutions.
- Delivered a parameter-driven executive dashboard surfacing a high-quality acquisition channel (4.8% conversion, 5x top traffic source) and a retention segment driving 24.9% of revenue from 13.9% of customers, backed by an executive recommendation memo.

Cloud Architecture for Bankruptcy Prediction [Repo] | *AWS (S3, Glue, Redshift, SageMaker)* — *September 2025 - October 2025*
- Designed and implemented a scalable cloud-native data pipeline on AWS to consolidate decentralized financial data and support real-time analytics.
- Built automated ETL workflows in AWS Glue to extract, transform, and load 3,000+ financial records into an AWS Redshift data warehouse, cutting manual processing time by 80%.
- Drove proof-of-concept development by training, optimizing, and deploying an ML risk prediction model in AWS SageMaker with 97% accuracy, leveraging cloud-native services to deliver measurable business impact.

Waste Classification with CNNs [Repo] | *Python, PyTorch* — *February 2025 - April 2025*
- Fine-tuned deep-learning (AI) models to classify 4,752 landfill images across 9 categories, managed through Agile development iterations.
- Achieved an 88.9% accuracy fine-tuning ResNet50, closely matching prior work using InceptionV3, leveraging significantly lighter and more efficient model architecture.
- Implemented feature engineering capabilities and data quality assurance processes for deep learning models, creating scalable distributed architectures in PyTorch for robust algorithmic deployment.
- Translated model results into clear visualizations and insights to showcase practical deployment in real-world sorting.

LEADERSHIP & INVOLVEMENT

Data Science Club (DotData) — Madison, WI
- Led team analysis of customer churn data in campus competition, developing a predictive model achieving 84% accuracy and placed 2nd among 15 teams.

Information Systems Society — Madison, WI
- Organized 8 speaker events with 50+ attendees, facilitating connections leading to job opportunities for members.

SKILLS & INTERESTS

Data & AI Engineering: Python (Pandas, PyTorch, BeautifulSoup4), SQL, R, HTML, ETL/ELT Pipeline Development, Cloud Data Architecture, Data Modeling, Feature Engineering, Scalable Distributed Architectures, Proof-of-Concept Development, Machine Learning, Predictive Modeling, Quality Assurance.

Cloud & BI: AWS (S3, Glue, Redshift, SageMaker), Google Cloud Platform (GCP), Microsoft Azure, PostgreSQL, Snowflake, BigQuery, Tableau, PowerBI

Data Governance, Quality & Management: Data Governance, Data Quality Assurance, Data Cleaning & Validation, Metadata Management, HIPAA Compliance Standards, Statistical Analysis, A/B Testing & Experimentation, Version Control (Git), Linux (Ubuntu, RedHat), Windows, macOS, Microsoft 365, Google Workspace.