Risshabh Bhardwaj

risshabh@gmail.com | (609) 373-3358 | www.linkedin.com/in/risshabh | https://github.com/RisshabhBhardwaj | Available Dec 2026

EDUCATION

Purdue University — **West Lafayette, IN**

Bachelor of Science in Computer Science with a Concentration in Machine Intelligence — December 2026

Bachelor of Science in Artificial Intelligence with a Minor in Mathematics and Philosophy — December 2026

Master of Science in Artificial Intelligence with a focus on Machine Learning — Expected December 2027

- Coursework: Machine Learning | Natural Language Processing | Data Engineering | Systems Programming | Probability | Statistics

PROFESSIONAL EXPERIENCE

Accenture — **New York City, NY**

Tech Architecture Summer Analyst — June 2026 – August 2026

- Architected and developed AI platforms, systems, and tools while contributing to strategy and delivery for some of the world's largest media, entertainment, cosmetics, athletics, and nonprofit organizations, from discovery through technical implementation
- Built and **deployed a 6-agent LLM system** on AWS that **compressed a multi-week creative campaign workflow into minutes**
- Developed a schema-harmonization pipeline that unified inconsistent datasets into a single canonical schema, **removing 98% of the manual reconciliation** process required from data scientists, **reducing a week-long process to approximately one hour**
- Implemented a RAG retrieval system for a nonprofit reaching **92%+ precision, recall, and F1,** reducing a multi-hour process to seconds

L'Oréal — **Clark, NJ**

Software Engineer Intern — June 2025 – August 2025

- Optimized time-consuming processes by implementing AI-driven automation pipelines, **saving an estimated $4M annually and boosting chemists' efficiency by 15%** while automating data processing and analysis for predictive insights and formula optimization
- Stability Testing Automation (Python) - Automated CV and LLM pipeline providing chemical stability decisions and quantifying formula separation with **99% accuracy on a sample set of 7,000 formulas, replacing manual visual inspection**
- Data Processing Pipeline (R) - End-to-end analytics pipeline that automated data processing and statistical analysis (correlation + clustering) while deploying AI models for predictive insights and formula optimization, saving L'Oréal data scientists significant time

Educational Testing Service (ETS) — **Princeton, NJ**

Machine Learning Engineer Intern — June 2024 – August 2024

- Worked as a full-stack developer creating advanced security features for exams such as GRE, TOEFL, SAT, and AP by developing and implementing advanced AI security features using Tensorflow now **running in production** within ETS's Test Security as a Service

PROJECTS

Memory-Constrained Forecasting Research | Python, Polars, MLX

- Researched what a forecasting system should keep resident when memory is the binding constraint, tested on electricity-price tail risk
- Established that compact learned vector state beats **14× more raw scalar history at identical byte cost**, significant in all 12 test cells
- Built on a **329M-row vintage-correct corpus** of 15-minute nodal prices **over a survivorship-safe 499-node universe**
- Preregistered kill gates before running anything; two of three hypotheses failed and were published as failures
- Evaluated under 99% CVaR and pinball loss with block bootstrap, Diebold–Mariano with HAC, and Holm correction hypothesis families
- Cut peak memory from roughly **10 GiB to 2.8 GiB with zero swap** and bit-exact outputs after root-causing three eager allocations

Gridiron Oracle - Causal Forecasting and Execution | Python/C++ | gridironoracle.vercel.app

- Live NFL projection system that forecasts per-player stat lines each week, with every feature constrained to data available before kickoff
- Beats naive and trailing-3 baselines on **74 of 75 stat-position cell-seasons** from 2021–25, roughly **7% better MAE**
- Runs signals through a lock-free C++ ring buffer measured at **0.46–0.50 µs p99 and 100M+ msg/s**, 17–23× better than mutex baseline
- Serving predictions **match out-of-fold artifacts exactly across 80,404 rows**, and no model promotes without passing a frozen gate
- Tracks latent player form with a causal Kalman state feeding an ensemble of GBDTs and transformers via per cell Ridge stacking

Mainstay - Multi-Tenant Revenue Operations Platform | TypeScript, Fastify, Postgres, Redis, BullMQ

- Multi-tenant platform providing visibility into fragmented data, including calls, POS, and ad events, into a single normalized data spine
- **Deployed live system with a pilot venue**, backfilling historical data from 2025 onward to generate nightly attribution and funnel marts

Autonomous Orbital Optimization and Satellite Constellation Design for Disaster Response | Python, XGBoost

- Replaces hours-long orbital simulation with a learned surrogate to search constellation designs against disaster-response revisit time
- Cuts modeled worst-case revisit gap over disaster regions from 48h to under 15h through Pareto optimization, **43.9% over baseline**
- Predicts revisit windows at **held-out R^2 0.892 against CV R^2 0.897 on 7,000+ simulations**, constrained inside physics bounds at **98.7%**

SKILLS

Programming Languages:

- Python | TypeScript | SQL | C++ | Java | C | JavaScript

ML & LLM Systems:

- PyTorch | TensorFlow | XGBoost/LightGBM/CatBoost | Scikit-learn | PyMC | RAG | Embeddings | Agent Orchestration

Backend & Data:

- Postgres | Redis | Celery | BullMQ | FastAPI | Fastify | Next.js | Docker | Alembic | Polars | DuckDB | Neo4j | AWS | GCP