Pradyot Bathuri

Bloomington, IN | pbathuri@iu.edu | (930) 904-5450 | github.com/pbathuri |
linkedin.com/in/pradyotbathuri

EDUCATION

Indiana University Bloomington *Anticipated graduation: 05/28*
Dual BS Computer Engineering and BS Mathematics, Program II
CGPA: 3.5

Relevant coursework: ENGR-E 517 High Performance Computing, MATH-M 343 Differential Equations I, MATH-M 311 Calculus III, MATH-M 303 Linear Algebra, STAT-S 350 Statistical Inference, PHYS-P 222 Physics II, ENGR-E 222 Intelligent Systems II.

PUBLICATIONS

Bathuri, P. "Training-Horizon Effects in LLM-Assisted Quantum Portfolio Optimization." *Under review, QNLP AI 2026.*

RESEARCH

qhpc_cache: Empirical L1 cache characterization of quantitative finance kernels *January 2025 to May 2026*
Independent undergraduate research
github.com/pbathuri/finance-cache-hpc

- Ran an empirical cache behavior study for four numerical finance kernels: Cholesky factorization, Monte Carlo simulation, GARCH estimation, and GEMM.
- Used PAPI hardware counters to measure access patterns, L1 miss rates, and bandwidth saturation under varying workload regimes.
- Wrote a technical report; results inform cache-aware reordering of inner loops in derivatives-pricing routines.

Quantum computing foundations for HPC and finance applications *April 2026 to present*
Undergraduate research assistant, Indiana University Luddy School
Advisor: Prof. Yuxi Hong

- Studying the linear algebra behind single and multi-qubit gates, tensor products, entanglement, and the measurement postulate.
- Attending weekly research meetings, writing progress reports, and co-authoring a structured technical reference in a shared LaTeX project.
- Focus on the matrix-algebra structure of quantum operations and its connection to large-scale sparse computation.

RegimeFactorZoo: Machine-learning factor models with a regime-stability extension *May 2026 to present*
Independent research
github.com/pbathuri/RegimeFactorZoo

- Building a reproducible factor-model and machine-learning pipeline on public data from the Kenneth French Data Library, the Open Source Asset Pricing dataset, and FRED.
- Replicating Fama-French three-factor and momentum models; extending to Lasso, Ridge, ElasticNet, gradient-boosted trees, and sparse-Bayesian factor models.
- Testing whether sparse factor selections persist across VIX-defined market regimes, building on the 2026 Baba-Yara working paper *In Search of Sparsity*.

TECHNICAL SKILLS

Programming languages: Python (2 yrs), C/C++ (2 yrs), Java (1 yr), R (1 yr), LaTeX.
Python scientific and ML stack: pandas, NumPy, SciPy, scikit-learn, XGBoost, PyMC, Matplotlib, Seaborn, Jupyter.
HPC and systems: CUDA, PAPI hardware counters, SLURM, OpenMP, MPI basics.
Other: Git, GitHub, RAG pipelines (LangGraph, Qdrant).

CERTIFICATIONS

IBM Generative AI Engineering Professional Certificate (Coursera).
Wharton Online (Coursera): Fundamentals of Quantitative Modeling.

LEADERSHIP

Head of Technology, IU Data Science Club *2026 to present*
Designing and leading biweekly hands-on workshops for a ten-member group on data engineering, machine learning, and statistical analysis. Skills built: technical writing, weekly progress presentations, and peer mentoring on applied data-science workflows.