Hritwik Biyani

Stony Brook, NY | biyanihritwik@gmail.com | (934) 227 – 8432 | linkedin.com/in/hritwikbiyani | hb-hello.github.io

EDUCATION

Stony Brook University — **Stony Brook, New York**
Master of Science in Computer Science (GPA: 4.0/4.0) — *August 2025 – August 2027*
Relevant Coursework: CSE 535 (Distributed Systems), CSE 526 (Operating Systems)

MIT World Peace University — **Pune, India**
Bachelor of Technology in Computer Science — *July 2017 – July 2021*

TECHNICAL SKILLS

Programming Languages: Java, Python, JavaScript, Scala, Linux, SQL, NoSQL, GraphQL, HTML, CSS/SASS, JSON
Tools / Frameworks: Git, Snowflake, Databricks, AWS, dbt, Node.js, Flask, Neo4j, Docker, Airflow, PostgreSQL, React
Skills: Software Development, API Design, Data Integration, CI/CD, Unit Testing, ETL, Data Structures, REST APIs

WORK EXPERIENCE

ZS Associates — **Pune, India**
Senior Data Engineer — *July 2023 – July 2025*

- Developed a FastAPI service implementing advanced prompt engineering with LLMs (Gemini/GPT-4) to deliver real-time, context-aware insights, streamlining complex research workflows for pharmaceutical KAMs.
- Engineered a Node.js batch-processing pipeline to automate data ingestion from 20+ sources (REST APIs, FTP, SQL), replacing manual form entry and reducing onboarding time by 2-3 weeks per source.
- Led a team of 4 to architect a scalable Data Change Request system using Azure Functions and SQL Stored Procedures to handle high data volumes, processing 4,000+ requests and saving ~40 hours of manual cleaning.
- Designed the data warehouse architecture and implemented data pipelines across various platforms such as Snowflake, Databricks, dbt and Python for multiple clients, accelerating complex analytics by ~50%

ZS Associates — **Pune, India**
Data Engineer — *November 2020 – June 2023*

- Built a Node.js backend for real-time contracting scenario analysis, processing tens of thousands of rows via configurable algorithms to optimize contract selection and boost ROI by >20%
- Developed a full-stack graph visualization tool using D3.js and Neo4j, writing optimized Cypher queries and implementing lazy loading to efficiently render complex health system hierarchies with 10,000+ nodes
- Engineered a scalable Databricks Spark pipeline to process 30TB+ of complex, nested JSON datasets, utilizing strategic partitioning to optimize compute costs and handle Transparency in Coverage data.

AnalyticsDomain — **Pune, India**
Web Development Intern — *February 2020 – May 2020*

- Designed and deployed Course and Event management applications with Python/Django, enhancing the user experience for 100+ users with JavaScript-driven features like interactive quizzes and charts.

PROJECTS

Practical Byzantine Fault Tolerant System | *Stony Brook University* — **August 2025 - October 2025**

- Developed a high-throughput PBFT consensus protocol in Java featuring a collector node pattern to reduce communication overhead along with Threshold Signatures and Checkpointing
- Validated system resilience by simulating complex Byzantine behaviors, including equivocation, timing attacks, and 'in-the-dark' scenarios, ensuring consensus integrity handling malicious nodes

Sharded Transactional Datastore (2PC + Paxos) | *Stony Brook University* — **October 2025 - December 2025**

- Built a scalable transaction engine (~5000 tps) in Java supporting Cross-Shard ACID transactions via Two-Phase Commit (2PC) over Paxos clusters, complete with re-sharding and cluster reconfiguration capabilities.