

# Dhruv H

8884293860 | [dhruvh0530@gmail.com](mailto:dhruvh0530@gmail.com) | [linkedin.com/in/dhruv-trivedi-3005d](https://www.linkedin.com/in/dhruv-trivedi-3005d) | [github.com/dhruv-0512](https://github.com/dhruv-0512) | [dhruvtrivedi.me](https://dhruvtrivedi.me)

## SUMMARY

---

Third year Information Science student at BMS College of Engineering interested in systems, machine learning, and performance focused software engineering. I enjoy understanding the mathematics and tradeoffs behind complex systems and turning that understanding into practical software.

## EDUCATION

---

### B.M.S College of Engineering

*Bachelor of Engineering in Information Science*

- CGPA: 9.232/10.0

Bengaluru, Karnataka

*Aug. 2023 – Present*

### Base PU College

*Pre-University Course (PUC)*

- Percentage: 92.16%

Bengaluru, Karnataka

*2021 – 2023*

### V.V.S Sardar Patel High School

*Secondary School Leaving Certificate (SSLC)*

- Percentage: 99.36%

Bengaluru, Karnataka

*2021*

## PROJECTS

---

**dvvara**: Malicious URL detection via Bloom Filter | *Python, FastAPI, Redis, PyPI* | [visit](#) ↗

- Published a PyPI package using a Bloom Filter (the same technique used by Chrome Safe Browsing) to detect malicious URLs at 1/100th the memory of a hash set, indexing 268K+ threats from URLhaus, PhishTank, OpenPhish, and CERT Polska into a 5.14MB filter sized for 3M URLs at 0.1% target FPR.
- Designed a two stage lookup (~0.003ms bloom check, 145K URLs/sec, PostgreSQL confirmation only on hits) validated to 0 false negatives and 0 false positives across 100K test URLs, and self-hosted the full stack (FastAPI, PostgreSQL, Redis) on a single AWS EC2 instance via Docker Compose.

**QueryMind**: Natural language to SQL platform | *FastAPI, PostgreSQL, Redis, Kafka, ChromaDB, Gemini, React*

- Built a text-to-SQL platform where users upload CSV/Excel files and query in plain English; RAG pipeline retrieves DDL embeddings from ChromaDB, Gemini 1.5 Flash generates validated SQL executed against isolated per-user PostgreSQL schemas with Redis result caching and JWT auth.
- Designed event-driven architecture with Kafka (KRaft) across 4 topics and an Audit Consumer persisting all events to PostgreSQL, with Redis sliding window rate limiting and read-only SQL validation enforced before every query.

**Exoplanet Atmospheric Retrieval** | *Python, XGBoost, Scikit-learn, Optuna* | *Manuscript under review, Scientific Reports*

- Owned model training and pipeline design on a team research project benchmarking MLP, SVR, and XGBoost on the Ariel Big Data Challenge dataset to retrieve 5 gas abundances and equilibrium temperature from transmission spectra; XGBoost generalised best to real JWST NIRSpec observations of WASP-39b, matching published retrieval results.
- Identified preprocessing as the dominant performance factor (52-bin wavelength interpolation, log-scaling, sample-wise normalisation), and validated generalisation on out-of-distribution real JWST observational data across all five atmospheric gases.

## TECHNICAL SKILLS

---

**Languages**: C, C++, Python, Java, JavaScript, HTML/CSS

**Frameworks/Libraries**: FastAPI, Express.js, Node.js, React, Spring Boot, PyTorch, LangChain, Pandas, NumPy

**Infrastructure & Cloud**: Redis, PostgreSQL, MySQL, MongoDB, ChromaDB, Docker, Docker Compose, Kafka, AWS

**Developer Tools**: Git, GitHub, VS Code, Linux, UNIX, Postman, Jupyter Notebook, Claude Code, Cursor

**Concepts**: Data Structures & Algorithms, Probabilistic Data Structures, Distributed Systems, OOP, REST APIs,

Microservices, Machine Learning, Deep Learning, Retrieval Augmented Generation (RAG), Vector Embeddings,

WebSockets, Competitive Programming (Leetcode/Hackerrank)

## CERTIFICATIONS

---

**Nutanix Certified Associate (NCA)**: Hyperconverged infrastructure and Nutanix AOS administration.

**JPMorgan Chase: Software Engineering Job Simulation (Forage)**: Kafka integrated Spring Boot microservice with JPA and REST.

**NVIDIA Generative AI Explained**: Fundamentals of generative AI, GANs, and transformers.