Raghav Doshi

Email : raghav.doshi@research.iiit.ac.in
linkedin.com/in/raghav-doshi-7a576a2b4
Mobile : +91 737-890-5816
github.com/ThePurpleJedi/

Education

International Institute of Information Technology, Hyderabad — Hyderabad, India
B.Tech. & M.S. by Research in Computer Science and Engineering — *Jul. 2023 – Present*

- CGPA: 8.87/10
- Dean's List, 2024-25 | Merit List, 2023-24

The Bishop's Co-Ed School and Junior College, Kalyaninagar — Pune, India
Indian School Certificate — *Jun. 2010 – May 2023*

- ISC Grade XII: 98% | ICSE Grade X: 98.2%
- All-India Rank 1281 (top 0.7%) – Joint Entrance Examination (Advanced) 2023 | Top 0.43% in JEE (Mains)

Research Experience

SyPy Security and Privacy Research Group — Hyderabad, India
Undergraduate Research Student — *Apr. 2025 – Present*

- Investigating various fields in the intersection of ML and cybersecurity such as adversarial ML and backdoor stealthiness, along with more fundamental topics such as linear mode connectivity. Advised by Prof. Ankit Gangwal
- Analysed and presented papers from top conferences to the group, discussing potential directions for future work such as improving attacks on LLM watermarks and detecting architectural backdoors in neural networks using topological evolutionary dynamics.

Major Projects

- **Aligning LLMs Explicitly for Ambiguity** | *PyTorch, Transformers, KTO* | November, 2025
  Designed and applied **information-theoretic frameworks** for quantifying ambiguity in user queries to motivate aligning LLMs explicitly to handle ambiguity by asking clarifying questions. Benchmarked **Kahneman-Tversky Optimization** (KTO) against Supervised Fine-Tuning (SFT) for generating clarifying questions, building upon established alignment research, demonstrating that KTO can help improve performance.

- **Firefly** | HackIIIT Winner 2025 | *LSP, MOYA Multi-Agent Framework* | March, 2025
  As part of a hackathon, architected a **multi-agent system** leveraging **Language Server Protocol** (LSP) and the MOYA framework for **automating** the generation of technical documentation and synthesising code given simple example usages and documentation using specialised agents.

- **Sparse MoE Transformer with Custom GQA** | *PyTorch, Hugging Face, MoE, GQA* | October, 2025
  Exploring the internal mechanisms of modern transformer variants: implemented and trained a **sparse Mixture-of-Experts transformer** from scratch with a **custom GQA implementation**, two routing algorithms (hash routing and token choice top-k routing) and a learnable load balancer. The GQA setup provided similar performance to standard MHA with upto **75% reduction in KV cache size** and lesser computational load overall.

- **Quantisation and Compression** | *PyTorch, Transformers, Quantisation, Bitsandbytes* | September, 2025
  Conducting feasibility analysis of edge-based LLM inference: applied whole-model and selective component **quantisation**, integrating bitsandbytes for 8-bit, 4-bit quantization, and NF4 nonlinear quantisation. Evaluated trade-offs in efficiency and accuracy, showing that NF4 gives the best memory efficiency among the given methods.

- **Mercurys** | *Java, Object-Oriented Programming, Multi-threading, Sockets* | February, 2022
  Understanding key concepts from OOP, OS, and networks: applied multiple **fundamental OOP concepts** and **clean code concepts** along with multi-threading and sockets to create a CLI chat application in Java capable of transferring multiple file types without external libraries.

Teaching Roles

Teaching Assistant
International Institute of Information Technology, Hyderabad

- **Introduction to Information Security** | Spring 2026
  Designing and delivering lectures to a class of 100 students on the fundamentals of ML and ML security; designing and evaluating both written assessments and paper presentations
- **Computer Systems** | Spring 2026
  Designing questions for students and conducting weekly live tutorials
- **Introduction to Computer Organisation** | Monsoon 2025
  Designed questions for assignments and quizzes
- **Fundamentals of Computer Organisation** | Monsoon 2025
  Edited lecture videos and designed questions for assignments and quizzes

Activities

Theory Group, IIITH Hyderabad, India
Club Coordinator *Jun. 2025 – Present*

- Spearheaded the launch of new academic competitions including the Integration Bee, attracting over a hundred participants

Moderator *Jun. 2024 – May 2025*

- Facilitated seminars and weekly discussions on theoretical topics from fields including algorithms, graph theory, information and coding theory, quantum information and computing, and the mathematical foundations of computer science

IIITH Society for Applied Quantum Computing Hyderabad, India
Club Member *Jun. 2025 – Jan. 2026*

- Delivered a talk motivating the quantum computing paradigm at the Qiskit Fall Fest 2025 (IIIT Hyderabad), held in collaboration with IBM Quantum
- Presented foundational theoretical topics in the field of quantum computing such as the No Cloning Theorem and Bell's inequality

Achievements

1st Place, DeccanCTF Infinium: Ranked 1st among 400+ participants, showcasing proficiency in code-breaking, pattern-matching, network security, and creative approach to problems
1st Place, HackIIIT 2025: Won top prize for Firefly, an LSP-based multi-agent system for automated documentation/code generation
3rd Place, DeccanCTF 2025: Ranked 3rd among 350+ participants, demonstrating capabilities in OSINT, cryptography, and problem-solving
Best Implementation, HackIIIT 2024: Recognised for building a clean, usable campus-wide courier-tracking portal

Technical Skills

Programming/Scripting Languages
Python, C (Advanced) | Java, C++, JavaScript (Proficient) | PHP, Racket, Lean, Bash (Familiar)
Libraries
PyTorch, Transformers, Bitsandbytes, Accelarate, TRL, SciKit, Pandas, NumPy, SciPy
Frameworks & Technologies
CUDA, MediaPipe, NodeJS, React, NextJS, MySQL, ExpressJS, Flask

Selected Coursework

Advanced Natural Language Processing | Principles of Information Security | Introduction to Information Security | Quantum Error-Correcting Codes | Advanced Computer Networks | Software Programming for Performance | Probability & Statistics | Algorithm Analysis and Design | Automata Theory | Data and Applications