

Richi Dubey

Phone number: [4703387465](tel:4703387465) | Email: richidubey@gmail.com |
Linkedin: <https://www.linkedin.com/in/richidubey/> | GitHub: <https://github.com/richidubey> |
Blog: <https://rtemswithrichi.wordpress.com/>

EDUCATION

Georgia Institute of Technology – Atlanta, GA *August 2024 - May 2026*
Master of Science, Computer Science (Machine Learning Specialization), GPA 4/4
Relevant coursework: Computer Vision, Natural Language, Deep Learning, Graduate OS (Teaching Assistant)

Birla Institute of Technology & Science, Pilani – Goa, India *Aug 2017 - June 2021*
B.E.(Hons), Computer Science, GPA: 3.9/4
Relevant coursework: Machine Learning, Data Structures & Algorithms, Operating Systems (A grade)

EXPERIENCE

CERN (European Organization for Nuclear Research) – Geneva, Switzerland *October 2022 – July 2024*
Software Engineer

- [Developed](#) multi-threaded device drivers in C++ for a distributed [SCADA](#) system called [REMUS](#) that interfaces **1000+** diverse sensors deployed in CERN's accelerator areas. Also developed fault-tolerant networking programs for sensors, reducing downtime to less than **.001%**.

Oracle – Bangalore, India *July 2021 – September 2022*
Software Engineer

- Implemented **4+** features in a cloud application with microservice architecture using Java Spring Boot.
- Wrote terraform code to deploy Kubernetes infrastructure for the application on Oracle cloud. Deployed these codes across **50+** OCI data centres worldwide and reduced developer involvement by over **30%**.

OPEN SOURCE CONTRIBUTIONS

RTEMS – Sponsored by [Google Summer of Code](#) *Summer 2020*

- [Contributed 1000+ lines of code](#) to [RTEMS](#) for implementing the Strong Arbitrary Processor Affinity (APA) scheduler that allows dynamically relocating higher-priority tasks among processors and improves system schedulability by over **20%**. Published a [paper](#) and wrote a [blog](#) on the implementation.

RECENT PROJECTS

Improving LLMs performance on mathematical operations *Fall 2024*

- [Integrated](#) ALUs in transformer architecture to enhance the performance of LLMs. Training on GPT-2 with 10k epochs gives **0.002 average error** across operations, outperforming GPT for up to 1 decimal place.

Improving Masked Autoencoders' (MAEs) performance by Enhancing Data Augmentation with Generative Models *Fall 2024*

- [Finetuned](#) Diffusion model and CycleGAN models to improve performance (**over 9% decrease in loss**) of MAEs that use CNN and pixel wise loss to learn latent representation for improved downstream tasks.

Homography Projection for Surfaces in Art Paintings *Fall 2024*

- [Built](#) a homography projection app to visualise surfaces in an elevation view without perspective distortion.

PUBLICATIONS

R. Dubey, V. Banerjee, S. Hounsino, G. Bloom, *Strong APA scheduling in a real-time operating system: work-in-progress*, International Conference on Embedded Software (EMSOFT), 2021. [\[DOI\]](#), [\[Talk\]](#), [\[Poster\]](#)

SKILLS

Programming Languages: C/C++/C++14, Python, Java, SQL, Javascript

Tools/Frameworks: PyTorch, Scikit-Learn, Spring Boot, Kafka, Docker, Kubernetes, Git, Jenkins, Linux