Rohan Kumar

roku@bu.edu | 860-929-9772 | linkedin.com/in/rohan-kumar-in/ | Boston, MA

EDUCATION

Boston University, College of Engineering Expected May 2025
Bachelors of Science, Computer Engineering *Boston, MA*

- 3.95/4.0 GPA; Dean's List all semesters
- Kappa Theta Pi, IEEE-HKN, awarded Red Hat Collaboratory Undergraduate Research Award
- Relevant Coursework: Machine Learning, Deep Learning, Advanced Algorithms, Operating Systems

EXPERIENCE

Reltio May 2024 – August 2024
Machine Learning Engineering Intern *Remote*

- Expanded schema description generator using LLMs by rebuilding data pipeline and creating evaluation frameworks for LLM responses to fine-tune model and prompts
- Achieved 3.9x speedup of LLM calls by implementing recursive batching of prompts
- Improved performance of a Mistral-7b-instruct-v0.2 LLM by quantizing model and removing inefficient components, to attain a 30% reduction in GPU memory space utilization
- Engaged in weekly sprint meetings, submitted and completed Jira tickets, and revised documentation

Sloan School of Management, Massachusetts Institute of Technology May 2023 – Present
Machine Learning Research Intern *Boston, MA*

- Performing operations research under Prof. Dimitris Bertsimas, applying ML/AI and optimization in regards to AI in healthcare across 3 concurrent projects
- Redesigned data pipeline, improved accuracy of models (XGBoost, ANN, RF, Optimal Classification Trees) by over 10% in predicting patient readmission into a hospital system within 30 days of discharge
- Leveraged mixed-integer programming to optimize schedules for 10 physicians over fiscal quarters

PeacLab, Boston University College of Engineering September 2023 – May 2024
Machine Learning Research Intern *Boston, MA*

- Conducted AI/ML research for cloud systems under Prof. Ayse Coskun, AI4CloudOps RedHat project
- Administered experiments regarding cluster deployment of PraxiPaaS, an AI tool for Python package software discovery in the cloud, discovering 500+ packages at a time

PUBLICATIONS

[Accepted] Zhang, Z, Korver, L, Li, J, **Kumar, R** et al. "PraxiPaaS: A Decomposable Machine Learning System for Efficient Container Package Discovery". IEEE Int'l Conference on Cloud Engineering (2024).

PROJECTS

RGBD VSLAM w/ Object Detection | *MATLAB, Javascript, Java, Git* February 2024 – May 2024

- Led team of 5 engineers in accomplishing 3D object detection using YOLOv8 neural network on images processed by visual simultaneous localization and mapping (VSLAM)
- Configured VSLAM in MATLAB to capture important features in frames of video
- Employed YOLOv8 model to capture objects within each frame, returning corresponding features
- Implemented linear algebraic algorithm to project 2D features with camera pose into 3D space

SKILLS

Languages: Python, MATLAB, SQL, C/C++, Java, JavaScript, Bash, PowerShell, Verilog
Developer Tools: git, linux, conda, Docker, Kubernetes, Prometheus, Databricks, Jira, HuggingFace
Libraries: PyTorch, XGBoost, Transformers, Scikit-Learn, Gurobi, Numpy, Pandas, Matplotlib, Scipy, Jupyter