

Why Rohit?

I am eager to join your team in a **full-time Software Development / Machine-Learning / Data Engineering role**, bringing a proven record of shipping production-grade systems across cloud platforms, data-intensive back-ends, and ML workloads. Pursuing an M.S. (by Research) in Computer & Information Technology at Purdue University with a **4.0 GPA**, I thrive on converting complex technical challenges into elegant, high-impact solutions.

Experience at Amazon Web Services (AWS Bedrock)

During my internship at AWS Bedrock as a Software Development Engineer, I created **Hubble**, a centralized control plane that automates the entire lifecycle of Anthropic Claude models, eliminating 90 % of previously manual steps. I engineered and integrated **15 + APIs** that surface real-time health, capacity, and cost telemetry for thousands of replicas, giving teams the observability they lacked. By embedding automated capacity management and intelligent rebalance planning into the Forklift service, I cut inference latency and throttling rates in half and reduced mean time-to-resolution by **50 %**, materially improving the reliability of Bedrock's generative-AI offerings.

Experience at Razorpay

During my tenure at Razorpay as a Software Development Engineer, I led several impactful projects, including the re-architecture of the Harvester data service, which saved \$10-11K USD per month. I developed custom whitelisting logic for Maxwell and Debezium CDC collectors, securing around 90% of the company's data. Additionally, I migrated the SRM dashboard from TiDb to Pinot, reducing lag from five minutes to one, and significantly enhancing the experience for multiple teams and clients.

Machine learning Internship at Jaguar Land Rover

As a Machine Learning Intern at Jaguar Land Rover, I designed and implemented machine learning pipelines to extract critical information from handwritten sentences, achieving a 70% improvement over previous algorithms. I also developed exponential smoothing and LSTM-RNN models, boosting failure prediction accuracy from 80% to over 90%.

Research and Projects at Purdue University

At Purdue University, my research has focused on optimizing resource allocation for machine learning pipelines. I designed a metamodel for predicting CPU/GPU and memory requirements, which reduced costs by 25%. I deployed multiple LLM and CV pipelines on Kubeflow in a multi-node Kubernetes cluster, saving 40-50% of resources on average. Additionally, I built and maintained numerous data pipelines to support an NSF-funded cybersecurity education project.

Technical Skills and Contributions

My technical expertise encompasses a broad spectrum of programming languages and tools, including Python, C++, Java, TypeScript, and cloud computing platforms. I have practical experience with data pipelines, container management, and full-stack software development, as well as machine learning engineering and MLOps. My open-source contributions, such as enhancing security for the Trino gateway and developing a tool to convert LookML to Cube.js, demonstrate my commitment to collaborative development and continuous improvement.

Commitment to Excellence and Innovation

Across AWS, Razorpay, and Purdue, I have consistently delivered six-figure cost savings, high-performance back-ends, and ML systems that scale. I bring rigorous engineering, a collaborative mindset, and a passion for customer-centric innovation—qualities that align with the demands of a full-time **SDE/MLE/DE** position.

Thank you for considering my application. I look forward to discussing how my skills and drive can advance your organization's goals.

Sincerely,
Rohit Bankar

rohitbankar2000@gmail.com | +1 – 7657672732