

Deeksha Ramakrishna

deeksharamakrishna6@gmail.com | +1 240-351-0343 | [linkedin.com/in/deeksha-ramakrishna](https://www.linkedin.com/in/deeksha-ramakrishna) | Atlanta, GA

PROFESSIONAL SUMMARY

AI & Cloud Infrastructure Engineer with 4+ years of experience bridging ML Engineering, MLOps, and SRE — building and deploying production-grade ML systems on AWS. Proven impact: cut federal compliance check time by 90% via Agentic AI pipelines, achieved 92% anomaly detection precision, maintained 99.99% availability for provider-level hybrid cloud infrastructure and reduced MTTR by 25% across multi-cloud IaaS. Uniquely positioned at the intersection of AI/ML engineering and cloud reliability.

TECHNICAL SKILLS

Languages & Databases: Python, SQL, Bash, MongoDB, MySQL, MariaDB

AI/ML: Supervised & Unsupervised Learning, Deep Learning, Fine-Tuning, Agentic AI, RAG, LLMs, NLP, Generative AI, MLOps, Prompt Engineering, OCR

Frameworks & Tools: TensorFlow, Keras, PyTorch, Scikit-Learn, LangChain, LangGraph, HuggingFace, Pinecone, MLFlow, Spark, NLTK, OpenAI, Claude, Pandas, Matplotlib, Numpy,

Cloud, DevOps & SRE: AWS, GCP, OpenStack, Kubernetes, Docker, Terraform, CloudFormation, Ansible, Helm, Argo CD, Linux Administration and Troubleshooting, Prometheus, Grafana, Graylog, Quobyte, CI/CD, GitHub Actions, Jenkins, GitOps

PROFESSIONAL EXPERIENCE

AI and Cloud Infrastructure Engineer, Elef Capital (Algorithmic Trading Startup), Atlanta, GA Jan 2026 - Present

- Designed a low-latency options trading infrastructure using Terraform, Docker, and Kubernetes for zero-downtime execution in AWS.
- Engineered an end-to-end ML-driven trading strategy and sentiment analysis pipeline, deployed to AWS.
- Developed Python and Bash automation scripts for data ingestion and service releases to run strategy logic.
- Maintaining optimal infrastructure health via Linux administration, performance tuning, and Prometheus/Grafana observability pipelines.
- Maintaining reliable CI/CD pipelines using GitHub Actions for seamless IaC updates and service rollouts.

Machine Learning Engineer | Teaching Assistant (TA), UMD, College Park, MD Feb 2024 - Dec 2025

- Engineered an automated research proposal review pipeline using Agentic AI, RAG, and OCR (AWS Textract, Bedrock), reducing federal compliance check time by 90% and mitigating the risk of funding cancellations.
- Built a production-grade NLP translation pipeline, achieving 97% accuracy to enhance cross-lingual research; delivered proactive cloud support eliminating critical system downtime for stakeholders.
- As a TA, I directed technical training for 60+ engineers on Docker and Kubernetes orchestration.

Site Reliability Engineer, SysEleven GmbH, Berlin, Germany Jul 2021 - Dec 2023

- Maintained a hybrid multi-cloud IaaS (OpenStack, AWS S3) using Ansible and Terraform, ensuring 99.99% availability for enterprise tenants; configured and managed Linux-based KVM/QEMU hypervisors.
- Provisioned and maintained bare-metal servers with OS installations, OpenStack configurations (Nova, Cinder, Neutron), system patching, OS upgrades, and capacity planning.
- Pioneered the company's first ML initiative: designed an end-to-end automated anomaly detection system for platform logs achieving 92% precision over baseline methods.
- Expanded provider-level observability via Prometheus queries and Grafana dashboards/alerts; tracked SLIs/SLOs, reducing MTTR by 25%.
- Troubleshooting platform-level network connectivity and DNS failures across tenant environments using standard Linux diagnostic tools.
- Participated in on-call rotations and blameless post-mortems; wrote Python/Bash automation for operational tooling and maintained platform-wide CI/CD pipelines to reduce toil.

EDUCATION

Master of Engineering in Cloud Engineering — University of Maryland, College Park, MD, USA

Master of Science in Data & Knowledge Engineering (Applied ML) — Otto Von Guericke University, Germany

Bachelor of Engineering in Computer Science and Engineering — Dr. Ambedkar Institute of Technology, India

PROJECTS

- Clinical Decision System for detecting Cervical Intraepithelial Neoplasia using Machine Learning and Deep Learning which resulted in 95% balanced accuracy, 90% specificity, and 87% sensitivity.
- Designing an end to end deployment of the Opentelemetry application using Docker, Kubernetes, EKS, AWS