Sarabesh Neelamegham Ravindranath

sarabeshnr@gmail.com — +1-930-333-4115 — linkedin.com/in/sarabesh — github.com/sarabesh

PROFESSIONAL SUMMARY

Analytical and results-driven **AI/ML Engineer** with 4+ years of experience across **DevOps**, **MLOps**, and applied **Machine Learning**. Experienced in building scalable ML pipelines, training and deploying deep learning models, and operationalizing AI systems for production. Skilled at bridging research and engineering through automation, model governance, and infrastructure optimization. Hands-on expertise in **Computer Vision**, **Reinforcement Learning**, and **Retrieval-Augmented Generation (RAG)**, with a strong focus on reproducibility, observability, and large-scale performance.

Work Authorization: Currently on **F-1 OPT** (EAD valid through June 2026). Eligible for 24-month STEM OPT extension. Requires H-1B sponsorship thereafter (Change of Status).

TECHNICAL SKILLS

Languages & Frameworks: Python, Go, C ++, Java, Groovy, Flask, Spring Boot
Core ML Libraries: PyTorch, TensorFlow, Scikit-learn, Hugging Face, LangChain, Open3D, TensorRT
AI Domains: Computer Vision (CNN, ViT, YOLOv5, Diffusion), NLP (BERT, GPT, LLaMA2, RAG), Reinforcement Learning (PPO, TRPO, VPG)
MLOps Tools: DVC, MLflow, SageMaker, Prefect, Argo CD, GitLab CI/CD, Airflow
Cloud Platforms: AWS (Lambda, ECS, EKS, RDS, S3, Step Functions, IAM, VPC)
DevOps & Infrastructure: Docker, Kubernetes, VMware Tanzu, Terraform, Helm, Jenkins, Concourse, GitHub Actions
Databases: MySQL, DynamoDB, Neo4j, Qdrant, PuppyDB
Certifications: AWS Solutions Architect (SAA), Certified Kubernetes Administrator (CKA), Certified Kubernetes Security Specialist (CKS), Hashicorp Terraform Associate

EXPERIENCE

Research Assistant — Indiana University School of Optometry, Bloomington, IN — Jan 2025–Present

- Collaborating with faculty on the **Walking Project**, performing 3D terrain reconstruction and complexity modeling using physical and geometric measures.
- Built automated analysis pipelines for point cloud data, integrating Open3D and Python scripts to accelerate reconstruction and terrain metric computation.
- Improved reconstruction efficiency by **25%** through optimized data preprocessing and surface meshing workflows.

Emerging Tech Intern, AI Specialist — Musco Lighting, Des Moines, IA — May 2024–Aug 2024

- Developed a modular **MLOps pipeline** using **DVC** and **GitLab CI** to automate dataset versioning, validation, and model retraining, improving iteration speed by **40%**.
- Implemented **YOLOv5** training and batch inference pipelines integrated with **MLflow** for experiment tracking and model performance monitoring.
- Designed automated dataset syncing between **NFS** and **S3** using triggers, reducing data latency by **80%**.
- Enhanced MLOps observability with MLflow metrics dashboards, providing model drift and data change alerts in production environments.

DevOps Engineer, MTS II — VMware, Bangalore, India — Feb 2022–Jun 2023

- Contributed to the **VMware Tanzu Application Platform (TAP)** Release team, automating multi-cloud installation and validation pipelines across **EKS**, **AKS**, and **TKGS**.
- Designed and implemented new test suites in **Go** for E2E validation, increasing overall test coverage from 20% to 75%.
- Integrated CI/CD workflows into existing Helm-based infrastructure, improving release reliability and reducing manual testing overhead.
- Collaborated with cross-functional teams to optimize Kubernetes deployments and streamline version control integration for automated testing.

Cloud Engineer — Presidio, Chennai, India — May 2019–Jan 2022

- Architected and deployed enterprise-grade CI/CD pipelines using **Jenkins** and **GitHub Actions**, implementing **Infrastructure as Code (IaC)** via **Terraform** and **CloudFormation**.
- Built scalable cloud solutions on **AWS** for clients including *ViacomCBS*, *StatsPerform*, and *WEX*, optimizing cost.
- Led migration of Jenkins workers to Kubernetes pods, improving resource efficiency and build times by **35%**.
- Developed internal DevOps libraries integrating Slack and Sumo Logic APIs for real-time alerting and monitoring, reducing mean time to recovery (MTTR).

PROJECTS

PuppyDB — Minimal Vector Database *May 2025–Present*
Developed a lightweight **vector database** using **LMDB** for persistent storage and custom **HNSW** indexing for approximate nearest-neighbor retrieval. Implemented metadata management, brute-force and graph search modes, and persistence for large-scale embeddings. GitHub

GridWorld-PPO — Path and Route Planning with RL *Jan 2025–May 2025*
Implemented reinforcement-learning agents using **PPO** in a 2D grid-world simulator with dynamic obstacles. Integrated reward-shaping and imitation learning to improve convergence, visualized real-time trajectories, and benchmarked policy stability across randomized maps. GitHub

Neural-Recon — Image Reconstruction from fMRI *Sep 2024–Dec 2024*
Reconstructed natural scenes from brain activity using **latent diffusion models** and **CLIP embeddings**. Mapped fMRI signals to latent representations aligned with semantic image features; replicated experiments using the Natural Scenes Dataset (NSD). GitHub

HybridRAG — Graph vs Vector RAG Benchmark *Aug 2024–Dec 2024*
Benchmarked **Graph RAG (Neo4j)** against **Vector RAG (LangChain + Ollama)** architectures for LLM retrieval tasks. Evaluated retrieval latency, contextual precision, and scalability of hybrid pipelines for large knowledge graphs. GitHub

Mono-Cam DeepVO — Monocular Visual Odometry *May 2024–Jul 2024*
Reimplemented **DeepVO** for monocular visual odometry, combining **CNN (FlowNet-based)** feature extraction with **LSTM** temporal modeling. Trained on the KITTI odometry dataset to predict 6-DoF camera motion, achieving stable trajectory estimation and reduced drift error. GitHub

Sentiment Analysis — Fine-Tuning, Deployment & Monitoring *Jan 2024–Apr 2024*
Simulated an end-to-end MLOps workflow for deploying and maintaining a sentiment analysis model in production. Fine-tuned **BERT** on the **Sentiment140 dataset**, wrapped the model with **FastAPI**, and deployed on **Kubernetes (Minikube)** using Docker. Implemented batch prediction jobs, **data drift detection**, and **Slack-based alerting** for model degradation. Added automated retraining workflows triggered via CronJobs. GitHub

EDUCATION

Master of Science in Data Science *Aug 2023–May 2025*
Indiana University Bloomington — Luddy School of Informatics, Computing, and Engineering
CGPA: 3.90/4.00
Key Coursework: Applied Machine Learning, Elements of AI, Exploratory Data Analysis, Computer Vision, Applying ML Techniques in NLP, Deep Learning Systems, Autonomous Robotics

Bachelor of Engineering in Computer Science and Engineering *Jul 2015–Apr 2019*
Velammal Engineering College, Anna University — Chennai, India
CGPA: 8.39/10.00