ABHINAYSAI KAMINENI

Arlington, VA | (571) 473-6156 | kamineniabhinaysai@gmail.com
linkedin.com/in/abhinaysai-kamineni | github.com/askmy-stack |

PROFESSIONAL SUMMARY

AI/ML Engineer specializing in computer vision, time-series modeling, and production MLOps. Built 3D detection pipelines achieving 0.948 mAP@50 on biomedical imagery, benchmarked 15+ neural architectures on 916 hours of clinical EEG, and shipped multi-cloud ML infrastructure processing terabyte-scale data daily at Jio Platforms.

TECHNICAL SKILLS

AI / Machine Learning: PyTorch, TensorFlow, Hugging Face, scikit-learn, Computer Vision, NLP, Time-Series, Transformers, CNNs, LSTM, Mamba, Object Detection (CenterNet, YOLO, Faster R-CNN), Transfer Learning, MLflow

MLOps & Deployment: Docker, Kubernetes, Apache Airflow, CI/CD, Jenkins, GitHub Actions, Azure DevOps, Model Versioning, Reproducible Pipelines, Prometheus, Grafana, ELK

Cloud & Data: AWS (EC2, S3, SageMaker, Lambda), Azure, GCP, Terraform, Databricks, BigQuery, Apache Spark, Spark Streaming, Kafka, ETL, PostgreSQL, MongoDB, Neo4j

Programming: Python (Pandas, NumPy, PyTorch), SQL, Bash, R, JavaScript; Linux, Git, Jupyter, Weights & Biases

PROFESSIONAL EXPERIENCE

Sales Operations Analyst, Follett Higher Education Group · *Washington, DC* **May 2025 – Present**

- Operate fulfillment and inventory data workflows during peak academic cycles across online and on-campus systems, maintaining process-level data integrity for hundreds of SKUs per cycle.
- Apply CourseTracks adoption analytics for demand planning, check-in accuracy, and inventory reconciliation across multiple enterprise data sources.
- Adhere to PCI-compliant processes and IAM-aligned access control standards aligned with audit and security requirements.

Data Operations Engineer, Jio Platforms Limited · *Navi Mumbai, India* **Jul 2023 – Jul 2024**

- Engineered CI/CD automation for ML model deployment across AWS, Azure, and GCP using Jenkins, Docker, and Kubernetes; reduced time-to-production by 85% and release failures by 60%.
- Deployed Kubernetes-based ML inference microservices for real-time and batch workloads, reducing model-serving downtime by 40% and enabling auto-scaling under high-velocity traffic.
- Built and maintained 100+ Apache Airflow DAGs with Vault-secured authentication for ML and data pipelines, achieving 99.9% reliability across multi-terabyte daily volume.
- Optimized Spark Streaming jobs for real-time feature pipelines feeding production ML models, cutting data latency by 40% under high-velocity workloads.
- Provisioned cloud infrastructure using Terraform across AWS and Azure, achieving 30% cost reduction through right-sizing and resource optimization.
- Integrated Prometheus and Grafana observability with automated alerting on model and pipeline health, reducing mean time to detect (MTTD) by 45%.

Data Analyst, PHN Technologies · *Pune, India* **Mar 2023 – Jun 2023**

- Optimized ETL pipelines in BigQuery, improving processing efficiency by 30% and cutting analytics latency by 40%.
- Implemented key-based authentication framework in Apache Airflow, securing 75% of critical pipelines while ensuring uninterrupted deployments.
- Built 10+ monitoring dashboards surfacing data-quality issues and pipeline drift signals enabling the team to catch regressions before they reached downstream consumers.

KEY PROJECTS

EEG Seizure Detection — *Multi-Architecture Benchmark on Pediatric EEG* **Jan 2026 – Apr 2026**

- Benchmarked 15+ neural architectures (LSTM variants, Mamba, Mixture-of-Experts, CNNs, Transformers, classical ML) on 916 hours of CHB-MIT pediatric EEG across 24 patients under strict subject-independent evaluation; best AUROC 0.740.
- Built EDF-to-tensor preprocessing pipeline using MNE library; sanitized corrupted PhysioNet timestamps; implemented weighted dataloader with BCEWithLogitsLoss to handle class imbalance.
- Containerized full PyTorch training pipeline with Docker for reproducible cross-environment experiments on GPU instances.

BYU Bacterial Flagellar Motors — *3D Object Detection in Cryo-Electron Tomography* **Jan 2025 – Apr 2025**

- Engineered automated training and evaluation pipelines for 3D electron-microscopy detection using CenterNet, YOLOv10, and Faster R-CNN; achieved state-of-the-art mAP@50 of 0.948 and precision of 1.00 on noisy small-object benchmarks.
- Containerized preprocessing and inference workflows with Docker on AWS GPUs, enabling reproducible cross-environment deployments with CLAHE augmentation, noise modeling, and transfer learning.

NASA Landslide Predictive Analysis — *Time-Series Modeling for Geospatial Risk* **Aug 2024 – Dec 2024**

- Automated ML training and deployment using GitHub Actions and Jenkins (70% fewer manual operations); provisioned scalable AWS compute via Terraform with infrastructure-as-code for version-controlled, reproducible analytics environments.

EDUCATION

The George Washington University — MS in Data Science (Global Leaders Award) · GPA: 3.77/4.0 **Aug 2024 – May 2026**

Savitribai Phule Pune University — BE in Computer Engineering, Honors in Data Science · GPA: 3.6/4.0 **Aug 2019 – Aug 2023**