Nikhil Mandayam Adyapak

+1(608)395-5104 | nikhiladyapak31@gmail.com | linkedin/in/nikhil-adyapak | nikhiladyapak.github.io | Google Scholar

Seeking 2026 Summer Internships/Fall Co-ops in ML, MLOps, Data Science Internships

Education

University of Wisconsin-Madison Sep 2025 - May 2027
Masters in Data Science GPA - 3.65/4.00
Coursework: Machine Learning, Foundation Models, Database Design & Implementation, Statistical Models, Statistical Methods, Statistical Inference for Data Science | **Teaching Assistantship**: CS564 Database Management Systems

PES University - Bengaluru, India Aug 2019 – May 2023
Bachelor of Technology in Computer Science & Engineering GPA - 3.78/4.00
Specialization: Machine Intelligence & Data Science | **Awards**: MRD Scholarship (Top 20% of 1100+ students)
Coursework: Algorithms, Data Structures, Data Science, Linear Algebra, Machine Learning, Big Data, Cloud Computing
Teaching Assistantship: CS203 Statistics for Data Science

Work Experience

Bosch Global Software Technologies (BGSW) Bengaluru, India
Senior Software Engineer – AI Systems for Advance Driver Assisted Systems (ADAS) Aug 2023 – Aug 2025

MLOps (ML in Operations) & Distributed Cloud Workflows

- Delivered **MLOps platform** for pedestrian detection (**Azure ML Studio** + **MLFlow**) with end-to-end **ML lifecycle**; enabled **multi-GPU training** (**3x** faster) and cut GPU costs by **60%**.
- Designed and scaled a **Ray Cluster** based distributed **ML training** system on **Azure Kubernetes Service (AKS)**, reducing retraining cycles from **4 weeks** to **1 week** (**75% faster**) and supporting **30+** teams.
- Deployed an on-prem **Ollama LLM** with a Python FastAPI wrapper, enabling private, team-wide LLM access (**150+ users**) with zero external data egress.

Large-Scale Image Retrieval & Dataset Generation for Autonomous Driving

- Developed a **60M+ image-search engine** using **CLIP embeddings** + **ElasticSearch**, enabling edge-case retreival and cutting reviewer triage from **4 hours** to **10 minutes** (**95% efficiency gain**).
- Engineered **scene-understanding** image retrieval pipeline for **fine-grained dataset** generation: scene graphs (**hugging face transformers** + **ensemble of foundation models**) improving corner-cases for complex driving scenarios.

Embedding Reuse & Cost Optimization

- Integrated **CLIP embedding** re-use with **80%** cross-version alignment, avoiding full re-indexing of legacy image vectors.
- Streamlined **ML pipeline** provisioning on **AKS** (**Terraform, Helm, GitHub Actions, Argo**) for **multi-tenant** ADAS teams; integrated Prometheus monitoring/alerting, and reduced experiment setup from **3 days** to **5 minutes** (**99%** efficiency gain) for **30+** teams.

Machine Learning Operations (MLOps) Intern Jan 2023 – May 2023 | Jun 2022 – Jul 2022

- Built on-premise **MLOps** pipeline (**DVC** + **MLflow**), reducing model setup time from **1 Week** to **1 Day**, **85%** faster.
- Accelerated **multi-GPU** training for YOLOv5 and Detectron2 on **Azure ML Studio**, reducing training time by **67%**.

Technical Skills

Programming languages: Python, C, C++, Java, SQL, R
ML & Statistics: PyTorch, TensorFlow, HuggingFace, Transformers, LLM, CNN, GAN, scikit-learn, OpenCV, matplotlib
MLOps & Workflows: Ray, MLFlow, DVC, Argo Workflows, Azure ML Studio, Prometheus, Grafana
Cloud: Azure, AWS, AKS, Docker, Terraform, Azure DevOps, GitHub Actions, Linux
Databases & Tools: ElasticSearch, PostgreSQL, MySQL, MongoDB, FastAPI, Flask, Streamlit, GitHub

Projects and Publications

[1] Code Runtime Complexity Prediction - ERCICA (Springer), 2023 | Python, TensorFlow, sklearn, Streamlit, NetworkX
Predicted Big-O runtime complexity of C/Java/Python code using static analysis and ML classification with BiLSTM over Abstract Syntax Trees graph embeddings on IBM CodeNet dataset (**Accuracy: 96%**). **[View Publication]**

[2] Sentinel-Edge AI - MadData Hackathon 2026 (Qualcomm Track) | Llama 3.2, Langchain, RAG, FastAPI
Selected as 1 of 10 teams to build an air-gapped **Edge AI** security auditor on Snapdragon X Elite NPU; engineered a RAG pipeline with ChromaDB to detect legal risks and code vulnerabilities without cloud egress. **[View Project]**

[3] MLOps POC pipeline for Pedestrian Detection - 2023 | Python, DVC, MLFlow, DAG, Streamlit, SQLite, PyTorch
Implemented on-premise MLOps starter template with DVC + Detectron2 for end-to-end ML lifecycle.**[View Project]**

Awards and Services

Top Achiever Award & **Best Presenter Award**: BGSW, 2024, Data Engineering Unit: Ranked **1st** among **150+** engineers
Volunteer: BGSW, 2023-2025, Trained people with disabilities (Youth4Jobs) and taught CS to underprivileged students.
Teaching Assistant: PES University, 2022, Mentored **120+** students in Data Science and served as a panelist in a hackathon.