

Dhanush Kumar Antharam

Charlotte, NC • Open to Relocation

[+1 \(704\) 930-3938](tel:+17049303938) | adhanushkumar99@gmail.com | <https://www.linkedin.com/in/a-dhanush-kumar> |

<https://github.com/Dhanush-adk> | <https://adhanushkumar.com>

Professional Summary

Software Engineer in building cloud-native, AI-driven, and data-intensive systems on GCP and AWS. Strong background in algorithms, distributed systems, backend engineering, agentic AI and RAG pipelines. Skilled in designing scalable microservices, deploying LLM-powered applications, and optimizing real-time data workflows. Passionate about building reliable, intelligent systems that scale.

Technical Skills

Languages & Frameworks

Python, R, SQL, Bash/Linux Shell, HTML/CSS, R Shiny, Pandas, NumPy, OpenCV

Cloud Platforms & Services

AWS (EC2, Lambda, EKS, RDS, API Gateway, SageMaker, S3, CloudWatch, CloudTrail, VPC, IAM), Google Cloud Platform (Vertex Search Engine, RAG Engine, BigQuery, Cloud Run, Container Registry, Cloud SQL, Cloud Build), shinyapps.io

Agentic AI, LLMs & RAG Systems

Langflow, LangChain, LangGraph, Langfuse, OpenAI, Gemini (Vertex AI), Anthropic Claude, Embedding pipelines, Vector stores, Document Chunking, RAG Pipelines, Multi-agent orchestration, Tool-calling, Prompt Engineering, LLM microservices

DevOps & Infrastructure

Docker, Kubernetes (EKS/GKE), Jenkins, GitHub Actions, CloudFormation, CI/CD, Git, Bitbucket, RESTful APIs, API Gateway, Agile/Scrum

Analytics & Visualization

A/B & multivariate test design, Jupyter Notebook, Tableau, Power BI, Tableau Prep, DAX, Matplotlib, Seaborn

ML & Deep Learning

Regression, Classification, Clustering, Recommendation Systems, CNN, RNN, LSTM, GRU, Transformers, Model Deployment & Monitoring

Data Engineering & Backend

PostgreSQL, MongoDB, ClickHouse, Google BigQuery, ETL Pipelines, Distributed Systems, Data ingestion from BigQuery/ClickHouse → dashboarding

System Design & Algorithms

Scalable microservices, event-driven architecture, distributed cache design, API gateways, workflow automation; Proficient in arrays, linked lists, trees, graphs, heaps, recursion, dynamic programming

Experience

Chekout AI

Jul 2025 – Present

Software Engineer – AI & Platform Engineering

New York, NY (Remote)

- Designed a **one-click agent deployment system** where users upload custom knowledge bases and prompts; the system automatically spins up a dedicated **AI agent**, indexes content, and serves real-time, context-aware responses.
- Built an end-to-end agentic AI platform that generates dedicated agents with automated ingestion, chunking, embeddings, vector indexing, and RAG-based retrieval.
- Designed an Agent Deployment Engine supporting document ingestion, embedding generation, vector indexing, and Cloud Run microservice spin-up.
- Developed production-grade RAG pipelines using Vertex AI Search, embeddings, and Langflow that deliver accurate, context-aware responses.
- Built ingestion and deduplication pipelines improving lead and chat analytics accuracy by **32%**.

Juspay Technologies Pvt Ltd

Jan 2022 – Jun 2023

Software Engineer - Data Intelligence

Bangalore, India

- Architected a real-time analytics platform reducing latency from 3 hours to 3 minutes, powering production recommendations.
- Built ETL pipelines ingesting 2M+ daily records into BigQuery; deployed on EKS using Docker-based workflows.
- Developed a Golang scheduler for Python/R microservices and automated CI/CD pipelines using Jenkins + ECR + EKS.
- Implemented anomaly detection models with 85% accuracy: optimized cloud usage saving 76% in infrastructure costs.

Deeploop Technologies

Software Engineer - ML Domain

Aug 2021 – Dec 2021

Hyderabad, India

- Built **Business Genie**, a real-time computer-vision system (YOLOv5 + Deep SORT) for employee/customer re-identification with **81% precision**, logging entry/exit metadata into PostgreSQL for analytics.
- Deployed the model as a **Django microservice on AWS Lambda + API Gateway**, reducing inference latency and compute cost by **40%** while enabling seamless serverless scaling.
- Designed a **heatmap/hotspot detection module** (OpenCV, NumPy) to identify high-engagement zones and help optimize store layouts using time-series footfall analytics.
- Developed a **CNN-based skin-cancer detection model** (PyTorch on SageMaker Studio) achieving **83% ROC-AUC** with drift monitoring and deployed it as a managed SageMaker endpoint.

Projects

Estimating Vehicle Speed on Roadways Using RNNs and Transformers: A Video-based Approach

- Built a real-time pipeline (OpenCV + LSTM, GRU, Transformers) that predicts vehicle speed from surveillance video, LSTM model reached **94.25% accuracy / 3.96 RMSE on VS13** and **78.62% / 10.99 RMSE on I24-3D**, outperforming radar-free baselines.
- Optimized Transformer variant to handle long-sequence dependencies, sustaining **90% accuracy** at 30 fps with sub-200ms inference per frame.
- Packaged as a **camera-agnostic, non-intrusive module** for traffic-monitoring systems—no extra roadside sensors required—enabling scalable, data-driven road-safety analytics.
- Tech stack: Python 3, PyTorch, OpenCV, NumPy, Pandas, Matplotlib, CUDA-enabled GPU acceleration, Git / GitHub.

Education

University of North Carolina, Charlotte

Master of Science in Computer Engineering

Aug 2023 – May 2025

GPA: 3.70 / 4.00

Internship

King Machine

Software Engineering Intern

Jan 2025 – May 2025

Charlotte, North Carolina

- Built an automated file-management platform in **Python** that deduplicates and archives machine-log data cutting manual cleanup to zero.
- Packaged as a **camera-agnostic, non-intrusive module** for traffic-monitoring systems—no extra roadside sensors required—enabling scalable, data-driven road-safety analytics.
- Created a **regex-driven batch parser + GUI** that merges log files, flags key metrics, and runs statistical analysis and delivered a report that replaced a **3-hour manual workflow** with a 30-second process.
- Integrated **Git bare-repo version control and notification hooks (email/Slack)** so engineers receive real-time status alerts and rollback capability, increasing workflow visibility and data integrity.

Facilities and Management Dept. UNC Charlotte

Software Engineering Intern, (Full Stack Student Developer)

Nov 2023 – Dec 2024

Charlotte, North Carolina

- Designed and developed a commissioning-management dashboard using R Shiny, MongoDB, and HTML/CSS/JavaScript, deployed on shinyapps.io to streamline workflows and reduce paperwork by 50%.
- Built real-time data integration pipelines for tracking equipment performance (boilers, chillers, AHUs) and visualized key operational metrics to support predictive maintenance.
- Led Agile development cycles—including sprint planning, daily standups, and stakeholder demos—improving iteration speed and commissioning accuracy.

Publication

Abnormal Activity Detection using Deep Learning, published in the Springer Nature LNNS series and presented at the 5th International Conference, World S4 2021.

Link to Publication: https://link.springer.com/chapter/10.1007/978-981-16-6309-3_63