Vaibhav R. Dandala

(984)-244-8318 | vrdandal@ncsu.edu | Linkedin: vaibhav-dandala | github: VaibhavD2006 | vaibhavdandala.com

Education

North Carolina State University — Raleigh, NC
B.S. Electrical Engineering & Computer Engineering (Dual Degree) — Expected Dec 2028
Dual Concentrations: Machine Learning & AI, Computer Systems Software — GPA: 3.80

Summary

AI/ML engineer focused on end-to-end production tools, agentic LLM/RAG pipelines, and edge ML inference. Co-published at ICMLA '25 with NASA JPL; deployed real-time edge ML for DoD IoT systems; built and shipped KeyMail, an AI-native CRM/Marketing platform with 100+ realtors. Experience in training and deploying RAG pipelines, military edge deployment, and building production AI-native SaaS platforms.

Experience

UNC Department of Radiation Oncology — Chapel Hill, NC
Lead Machine Learning Engineer — Apr 2026 – Present

- Achieved 93% accuracy on vertebrae segmentation and lesion detection by co-developing a CNN pipeline combining MedSAM3 for segmentation and YOLOv8 for detection, trained on 65,000+ CT images.
- Reduced radiologist data prep time by 35% across active clinical trials by building segmentation and preprocessing pipelines, directly accelerating oncology research throughput.

NASA Jet Propulsion Laboratory — Remote
AI/ML Researcher — Jun 2025 – Feb 2026

- Built RAG pipeline using PineconeDB with hybrid BM25 + dense embeddings; 80%+ PDS3 filter prediction accuracy across MSL archives.
- Processed and indexed 10,000+ MSL mission records across structured and unstructured PDS3 archives, enabling planetary scientists to query decades of Mars rover data through a natural language interface for the first time.

Defense Advanced Research Projects Agency (DARPA) — Remote
ML Engineer Intern — Apr 2025 – Jul 2025

- Deployed real-time edge ML on IoT for pose estimation and object detection in military zones; synthetic data + quantization pipeline for battlefield hardware eval.
- Engineered annotation workflows across 3,500+ synthetic and real-world battlefield frames, directly improving model reliability for soldier-deployed edge hardware operating under strict compute constraints.

HavenSmart LLC — Remote
AI Developer & Team Lead — Apr 2024 – Jun 2024

- Deployed internal RAG chatbot (LangChain, PineconeDB, Docker) handling 5,000+ monthly employee queries on deployment processes and workflows.
- Eliminated 60% of manual onboarding overhead for engineering hires by replacing live CI/CD and architecture walk-throughs with a vector-search pipeline over internal NetApp docs, cutting ramp-up time across the team.

Research & Publications

- **LUMIN (Co-Author, ICMLA '25)** — Agentic LLM multimodal NASA PDS query system; hybrid RAG workflow over MSL archives with BM25 + dense embeddings; co-authored with JPL Lead Data Scientist Dr. Emily Dunkel.

Projects

KeyMail | key-mail.co | *Next.js, React, Supabase, Gemini, Gmail API, LinkedIn API, X API, Chrome MV3*

- AI-native marketing platform for realtors; Agentic CRM, inbox, campaigns/sequences, and marketing automation (LinkedIn/X).
- Built Key Moments: pg_cron-scheduled milestone engine invoking Claude with pgvector RAG for personalized draft generation; approval-gated agentic writes with immutable per-event audit traces across the delivery workflow.
- 100+ realtors on the platform with Doorify MLS partnership

LeapGPT | *FastAPI, Next.js, MongoDB, Pinecone, OAuth, OpenAI*

- Startup-ideation agentic LLM platform attracting 150+ entrepreneurs; RAG workflow trained on 300+ consulting whitepapers with significantly higher retrieval accuracy than baseline.

Skills & Tools

- **ML/AI:** PyTorch, CNNs, YOLOv8, MedSAM3, Computer Vision, MCP, Quantization, Model Fine-tuning
- **LLM & RAG:** LangChain, Agentic Self-Loop Orchestration, Hybrid BM25 + Dense Retrieval, OpenAI, Gemini, PineconeDB, pgvector
- **Systems & Dev:** Python, JavaScript, FastAPI, Next.js, Docker, AWS, Supabase, MongoDB, PostgreSQL, Git, CI/CD