

Contact

6025869169 (Mobile)
kolakonu@gmail.com

www.linkedin.com/in/
kolapoolakonu (LinkedIn)

Top Skills

Retrieval-Augmented Generation (RAG)
GenAI MLOps
Autonomous & Multi-Agent Systems

Languages

English
Yoruba

Certifications

edX Verified Certificate for Python for Data Science
Data Science
OSHA 10
Remote Pilot License
edX Verified Certificate for Probability and Statistics in Data Science using Python

Publications

Degradation and Failure Modes of 26-Year-Old 200 kW Power Plant in a Hot-Dry Desert Climate
Degradation and reliability of 26 years old Solar One/SRP PV power plant, Phoenix, Arizona

Kolapo Olakonu

AI Engineer | Machine Learning, GenAI, Agentic AI, MLOps | Scalable AI Systems
Phoenix, Arizona, United States

Summary

I'm an AI Engineer and Machine Learning professional with strong depth across data science, statistical learning, large-scale ML systems, and generative AI, focused on building production-ready, scalable, and reliable AI systems.

I've completed rigorous training across:

AI & Machine Learning (UT Austin) — ML, deep learning, NLP, computer vision, GenAI, and model deployment

Data Science (UC San Diego MicroMasters) — probability & statistics, machine learning fundamentals, and big data analytics using Apache Spark

Agentic AI Engineering (Advanced Track) — LLM engineering, retrieval-augmented generation (RAG), autonomous agents, Model Context Protocol (MCP), and GenAI MLOps

My work sits at the intersection of statistical rigor and systems engineering. I'm comfortable framing ambiguous problems, validating assumptions with data, and building intelligent systems that move from experimentation into real production environments.

I bring hands-on experience in:

Machine learning and statistical modeling, including feature engineering, model evaluation, and bias-variance trade-offs

Large language model (LLM) systems, including fine-tuning (QLoRA), embeddings, vector databases, and RAG pipelines

Agentic AI architectures, designing tool-using, memory-enabled agents capable of multi-step reasoning and autonomous task execution

Scalable AI deployment, using API-based services, Docker, and MLOps practices to monitor latency, cost, quality, and reliability

I focus on building AI systems—not just models: systems that are explainable, observable, cost-aware, and aligned with real business and operational constraints.

Experience

Self-employed

Applied AI Engineer (Independent Projects)

February 2025 - Present (1 year 4 months)

Designing and building production-ready machine learning and generative AI systems, including LLM-based applications, retrieval-augmented generation (RAG), autonomous agents, and scalable AI services.

Focus areas include ML system design, agentic AI architectures, and GenAI MLOps.

SolarPTL

Project Engineer

May 2022 - December 2025 (3 years 8 months)

Tempe, Arizona, United States

Nautilus Solar Energy, LLC

Asset Manager

2018 - 2020 (2 years)

Phoenix, Arizona Area

Operations

Monitor NSE solar power plants daily across seven US States and in Ontario, Canada, comparing output to expected using AlsoEnergy PowerTrack and ACTIS

Coordinate and track all required repairs

Prepare all monthly reports for NSE management and financing partners.

Maintenance

Manage Operations and Maintenance (O&M) Agreements for NSE's generating portfolio.

Oversee monthly/quarterly/annual O&M activities, as required.

Develop tools, procedures and business relationships to help decrease NSE maintenance costs over time.

Ensure compliance with applicable warranties.

Maintain the O&M log.

Harmon Electric

Residential PV Systems Designer and BI Analyst

February 2016 - September 2017 (1 year 8 months)

Phoenix, Arizona Area

- Reviewed site auditor report and solar contracts to produce detailed CAD drawings of complete, permit-ready electrical and roof layout plan sets – Produce 1-Line, 3-Line Electrical Drawings, Array Diagrams and Site Plans, Layouts and Support Details of an average of 3 customer projects per day.
- Participated in design team planning meetings focusing on developing expertise in PV projects design and reporting using Salesforce CRM.

NRG Renew

Asset Optimization - Performance Engineer

2014 - 2015 (1 year)

Actively monitored 120 Distributed Generation and Utility Scale PV sites on monitoring platforms

Managed Planning and Scheduling for site repairs and conduct weekly meetings with 3rd party O&M providers

Logged operations daily activities and developed methodologies to improve efficient maintenance visits to PV sites

Reviewed daily all service reports of maintenance activities prepared and sent to NRG by field techs after services completion

Reviewed Independent Engineering reports from Luminate LLC, generated comments for NRG Renew management

Identified and trended system's performance in a proactive manner. Reduce the length of time energy availability issues existing on the sites by creating work orders promptly to 3rd party O&M providers

Involved in On Site utility plant supervision and conducted performance tests in collaboration with TUV Rheinland engineers

Sun Valley Solar Solutions

Commercial PV Systems Engineer

2013 - 2014 (1 year)

Maintained close communication within the sales team & project management

Maintained close communication with Balance of System sales managers & application engineers

Simulated solar energy production estimates for commercial systems using PVSYST

Performed site shading analysis

Performed array-inverter sizing, designed array layout drawings for the sales proposal package

Selected PV modules & Balance of System equipment on commercial projects

K & B Design

CAD Solar Site Designer

2013 - 2013 (less than a year)

Scottsdale, AZ

Produced CAD drawings for residential solar systems

Arizona State University

Research Assistant - Photovoltaic Reliability Laboratory

2011 - 2013 (2 years)

Mesa, AZ

- Performed research on degradation and reliability of SRP's 26 years old Solar One power plant, Phoenix, Arizona
- Collected, analyzed and reported field aged array data of mono-crystalline silicon modules for reliability and performance
- Performed accelerated testing of crystalline silicon solar modules under controlled laboratory conditions
- Lab instructor for new graduate research students on lab equipment and outdoor module testing
- Performed water vapor permeation experiments using MOCON water vapor transport machine

National Engineering and Technical Company

Piping Layout Engineer

2000 - 2010 (10 years)

Lagos, Nigeria

Supervised piping and mechanical equipment installations onshore, offshore, and fabrication in Sardinia, Italy 2003

- Involved in two major Front End Engineering Design projects as lead piping designer and senior piping engineer respectively

- Prepared 25% of hydro test packages on Exxon Mobil EPC 1A offshore gas compression platform installation and commissioning 2005
- Promoted to the position of piping layout engineer from piping layout designer in 2005
- Participated on teams for onshore and offshore of oil and gas process piping facilities site surveys
- Produced detailed piping studies and general arrangement drawings from Computer Aided Design (CAD) model, piping support drawings
- Training in Caesar II for pipe stress analysis, PDMS, AutoPlant for layout 3-D designs

Shell

Assistant Well Site Drilling Engineer

1998 - 1999 (1 year)

Port Harcourt, Nigeria

Worked with in Well Site Drilling Engineer in the drilling supervision of four oil well drilling operations on onshore and offshore drilling rigs

Education

Texas McCombs School of Business

Post Graduate Program Certificate, Artificial Intelligence and Machine

Learning: Business Applications · (February 2025 - October 2025)

Udemy | Advanced Professional Training (3-Course Track)

Certificate, Agentic AI Engineer Track (LLMs, Agents, MCP &

MLOps) · (November 2025 - February 2026)

UC San Diego

MicroMasters, Data Science · (December 2023 - June 2024)

Arizona State University

Masters of Science in Technology, Alternative Energy Technologies

Eclipse Consult

Certificate, Database Development and Analytics

AI Engineering & Agentic Systems Dossier

Kolapo Olakonu

This dossier summarizes the AI engineering, agentic AI, production deployment, automation, cloud architecture, AI-native coding, and multi-agent orchestration capabilities acquired through Ed Donner's professional AI engineering curriculum, combined with prior engineering, analytics, and technical leadership experience. This document is designed to support: LinkedIn profile optimization ATS-ready resume modernization AI Digital Twin and conversational AI knowledge bases Portfolio websites and professional branding Recruiter outreach and technical positioning AI consulting and business development opportunities

Professional Background

Kolapo Olakonu is an engineering and AI systems professional with experience spanning: Solar engineering and performance engineering Asset management and operational analytics Machine learning and AI engineering LLM systems and autonomous AI agents Cloud infrastructure and production AI deployment AI automation and workflow orchestration AI-native coding and agentic software engineering Photography, videography, and digital media production Entrepreneurship and AI-enabled business automation

Core AI Engineering Competencies

LLM Engineering and Prompt Engineering Retrieval-Augmented Generation (RAG) QLoRA Fine-Tuning and Open-Source LLM Training Agentic AI and Autonomous Multi-Agent Systems Model Context Protocol (MCP) Production AI Infrastructure and Cloud Deployment AI Workflow Automation with n8n AI Voice Agents and Conversational AI Agentic Coding and AI-native Software Engineering Cloud Infrastructure, Observability, and AI Guardrails Docker, Terraform, GitHub Actions, and CI/CD AWS Bedrock, SageMaker, Lambda, API Gateway, S3, CloudFront LangChain, LangGraph, CrewAI, AutoGen, OpenAI Agents SDK Claude Code, Codex, Cursor, Copilot, OpenRouter, MCP ecosystems

Course 1 — AI Engineer Core Track: LLM Engineering, RAG, QLoRA, Agents

This 8-week immersive AI engineering track focused on practical implementation of modern LLM systems, fine-tuning workflows, RAG pipelines, autonomous agents, and production-ready AI applications. Key capabilities acquired: Transformers, tokenization, context management, and LLM architecture Frontier and open-source LLM orchestration OpenAI APIs, Ollama, HuggingFace, LangChain, and Gradio Advanced RAG pipelines using embeddings and vector databases Evaluation systems using MRR, nDCG, and LLM-as-a-judge Fine-tuning with QLoRA, LoRA, HuggingFace TRL, and Weights & Biases Structured outputs with Pydantic and multi-agent orchestration Deployment of fine-tuned models and production AI systems

Projects Built — Course 1

AI-powered brochure generator that intelligently scraped and summarized company websites Multi-modal airline customer support assistant with function-calling and Gradio UI Meeting minutes and action-item extraction system using Whisper and open-source models Python-to-C++ optimization system achieving extreme performance acceleration Enterprise AI knowledge-worker using advanced RAG architecture Price prediction AI system using frontier and fine-tuned open-source models QLoRA fine-tuned open-source LLM competing with frontier models Autonomous multi-agent deal-finding and notification platform

Course 2 — AI Engineer Agentic Track: The Complete Agent & MCP Course

This 6-week advanced course focused on autonomous AI systems, agentic orchestration, distributed agents, multi-agent collaboration, and MCP-powered workflows. Major frameworks and ecosystems: OpenAI Agents SDK CrewAI LangGraph Microsoft AutoGen MCP (Model Context Protocol) Async Python and distributed runtime systems Structured outputs and agent guardrails

Projects Built — Course 2

Career Digital Twin AI agent for employer interaction and professional representation AI SDR sales outreach automation agents with email workflows Deep Research multi-agent system with autonomous web research pipelines CrewAI-based stock-picker and financial analysis platform Collaborative AI engineering teams using CrewAI and coding agents Browser automation sidekick agent using LangGraph and Playwright Distributed AutoGen agent communication systems with gRPC runtimes Autonomous multi-agent trading floor powered by MCP servers and multiple tools

Course 3 — AI Engineer Production Track: Deploy LLMs & Agents at Scale

This production AI engineering course focused on enterprise deployment, cloud-native infrastructure, observability, scalability, security, and production AI operations. Core production technologies: Vercel, AWS, Azure, and Google Cloud Platform Docker containerization Terraform Infrastructure as Code GitHub Actions CI/CD pipelines AWS Lambda, Bedrock, API Gateway, CloudFront, App Runner SageMaker and Bedrock inference systems Aurora Serverless and cloud-native AI architectures LangFuse observability and production monitoring

Projects Built — Course 3

Production healthcare AI SaaS platform using Next.js and FastAPI Digital Twin Mk II production AI assistant deployed on AWS infrastructure Cybersecurity analyst agent using MCP and multi-cloud deployment Research agent with vector ingestion pipelines and S3 vector architecture Production-grade multi-agent financial planning platform AWS Bedrock-powered agent systems with observability and guardrails Terraform-managed multi-environment cloud infrastructure deployments CI/CD automation pipelines for production AI systems

Course 4 — AI Builder: Create Agents, Voice Agents & Automations in n8n

This course focused on no-code and low-code AI workflow automation, voice AI systems, business automations, RAG pipelines, and MCP-powered workflows using n8n. Major technologies: n8n workflow orchestration OpenAI, Gemini, DeepSeek, and OpenRouter integrations Slack, Gmail, Telegram, Google Sheets, Google Drive ElevenLabs voice AI and Twilio integrations Supabase vector databases and RAG systems Self-hosted n8n with Docker and Ollama MCP integrations and automation systems

Projects Built — Course 4

Autonomous financial portfolio rebalancing agent using Google Sheets and OpenAI Slack-integrated AI automation bots and workflow systems AI-powered Gmail drafting and notification systems Conversational AI voice agents using ElevenLabs and Twilio RAG-powered voice AI assistant with Supabase vector storage Self-hosted n8n AI workflows integrated with Ollama and DeepSeek Multi-agent lead generation and CRM automation system using MCP AI-powered sales automation and demo-booking workflows

Course 5 — AI Leader: Generative AI & Agentic AI for Leaders & Founders

This executive-oriented course focused on AI strategy, organizational transformation, AI project leadership, implementation planning, and enterprise AI adoption. Key topics: Generative AI strategy frameworks AI adoption and business transformation Executive AI decision-making AI project roadmap development Cross-functional AI team leadership RAG and fine-tuning strategy for enterprises Business ROI and AI implementation planning

Projects & Strategic Exercises — Course 5

Enterprise AI strategy planning exercises AI implementation roadmap frameworks Business use-case ideation for predictive, generative, and agentic AI Cross-functional AI decision-making simulations AI organizational structure and leadership planning

Course 6 — AI Coder: Complete Claude Code & Coding Agents Course

This advanced AI-native engineering course focused on coding agents, multi-agent orchestration, autonomous software delivery, and next-generation AI engineering workflows. Core technologies and ecosystems: Claude Code, Codex, Cursor, Copilot, Antigravity Vibe Coding and Agentic Engineering YOLO workflows, Ralph Loops, and autonomous coding patterns AGENTS.md and CLAUDE.md context engineering MCP servers, hooks, plugins, skills, and slash commands Sub-agents, swarms, and orchestrator systems Cloud sandboxes and remote execution systems Jira and GitHub-integrated autonomous development workflows

Projects Built — Course 6

AI-powered personal website with conversational Digital Twin Kanban project management platform integrated with AI chat interaction AI legal SaaS application for drafting legal agreements and downloadable PDFs Full-stack Next.js and FastAPI SaaS applications using coding agents AI trading workstation with real-time market data and strategy assistant Multi-agent software engineering swarms using Claude Agent Teams and Gastown Jira-to-GitHub autonomous coding workflows using MCP integrations Remote cloud-sandbox coding systems with Claude Code and Codex orchestration

Technical Ecosystem & Platforms

Programming & Development: Python FastAPI React & Next.js Docker Git & GitHub VS Code & Cursor IDE AI Frameworks & Agent Platforms: LangChain LangGraph CrewAI AutoGen OpenAI Agents SDK Claude Code Gradio Transformers PyTorch Cloud & DevOps: AWS Azure Google Cloud Platform Terraform GitHub Actions CI/CD pipelines AI Models & Platforms: OpenAI GPT models Claude Gemini DeepSeek Groq OpenRouter Ollama HuggingFace

Strategic Positioning

The combined training and project portfolio positions Kolapo Olakonu as an: **Applied AI Engineer & Agentic Systems Builder** with practical expertise spanning: LLM engineering and fine-tuning Production AI deployment and cloud infrastructure Autonomous AI agents and multi-agent orchestration RAG systems and vector architectures AI-native coding and coding-agent ecosystems Workflow automation and voice AI systems Production observability, monitoring, and AI guardrails Enterprise AI transformation and deployment strategy This portfolio aligns strongly with the rapidly emerging future of: Autonomous AI systems AI-native software engineering Human-directed coding-agent orchestration Production-grade enterprise AI Multi-agent collaborative systems AI-powered operational automation