

MARWAN ALJIAKLI

AI/ML Engineer | CTO | Computer Vision | LLM & RAG Systems | IoT & Embedded Systems
+966 572 221 939 | marwan2004000@gmail.com | marwan-aljiaikli.com | github.com/MarwanAljiaikli
linkedin.com/in/marwan-aljiaikli-7ba965241 | Jeddah, Saudi Arabia | Syrian National — Valid Saudi Iqama
Available immediately | Open to remote roles globally and KSA on-site | Zero visa friction for KSA employers

PROFESSIONAL SUMMARY

AI/ML Engineer and CTO who shipped 3 production AI systems — including rPPG (camera-based vital sign detection) and real-time industrial computer vision safety — while completing a B.Sc. in Artificial Intelligence with a 4.35 GPA. Deployed multi-node IoT ecosystems processing 100K+ sensor events/day at sub-50ms latency and LLM/RAG pipelines serving live health-tech users. National competition finalist (Future Minerals Pioneers). SDAIA- and KAUST-trained. Targeting Senior ML Engineer, Computer Vision Engineer, and Technical Lead roles globally.

TECHNICAL SKILLS

AI/ML & Deep Learning: PyTorch, TensorFlow, Keras, scikit-learn, supervised & unsupervised learning, transfer learning, model fine-tuning, hyperparameter optimization, MLOps, model evaluation & validation, machine learning pipelines

LLMs & Natural Language Processing (NLP): LangChain, RAG (Retrieval-Augmented Generation), prompt engineering, Transformer fine-tuning, embeddings, vector search, semantic search, GPT-4, Claude, Llama, Qwen, Hugging Face

Computer Vision: OpenCV, MediaPipe, YOLO (v5/v8/v11), rPPG (camera-based vital sign extraction), object detection & segmentation, real-time video analytics, RTSP streaming, edge inference, Edge AI

MLOps & Infrastructure: Docker, FastAPI, microservices, CI/CD, GitHub Actions, AWS IoT Core, Google Cloud, Linux server administration, REST API design, Node.js

IoT & Embedded Systems: ESP32, STM32, nRF52840, Raspberry Pi, Arduino, FreeRTOS (real-time embedded OS), bare-metal C/C++, MQTT, BLE 5.0, Modbus, I2C/SPI, OTA firmware updates, PCB bring-up

Data Engineering & Databases: PostgreSQL, FAISS, ChromaDB, InfluxDB, ETL pipelines, data warehousing, SQL data modeling, vector databases

Frontend & APIs: React.js, TypeScript, JavaScript, Chart.js, Recharts, Tailwind CSS, REST API integration, real-time dashboards

Programming Languages: Python, SQL, C/C++, JavaScript, TypeScript, Bash

PROFESSIONAL EXPERIENCE

Chief Technology Officer | BOHIO | Jeddah, Saudi Arabia

January 2025 — Present

- Defined and executed the full technology strategy and AI product roadmap for an early-stage AI startup, translating investor-level business goals into quarterly engineering milestones across 3 concurrent product lines.
- Architected all LLM integrations, RAG pipeline designs, and prompt engineering frameworks, enabling rapid AI feature iteration with zero rework cycles across 2 production products.
- Built cloud-native infrastructure (FastAPI microservices, Docker, CI/CD) supporting ~800 active users at 99.7% uptime, enforcing automated testing coverage and async-first engineering standards.
- Reduced time-to-validation for new AI features to 24–72 hours using LLM-assisted rapid prototyping, eliminating full-sprint resource commitment before proof-of-concept validation.
- Hired and scaled the engineering team from 1 to 4 engineers, negotiated 3 vendor contracts achieving ~22% cost reduction, and drove cross-functional alignment across product, engineering, and executive stakeholders.

AI Engineer | VLEED | Jeddah, Saudi Arabia

June 2025 — Present

- Engineered a production-grade rPPG (camera-based vital sign detection) engine in Python (OpenCV, MediaPipe) achieving 91% HR accuracy (± 2.5 BPM MAE) vs. clinical pulse oximeter baseline at 28 FPS on live facial video — deployed to mobile SDK serving 300+ users.
- Architected and deployed scalable RAG pipelines (LangChain, FAISS, ChromaDB) on Dockerized FastAPI microservices, delivering AI-driven health insights with <140ms average query latency in production.
- Consolidated wearable device APIs and clinical lab data into a centralized SQL data warehouse, reducing downstream ML model preparation time by 65% and enabling 4 models to share a single source of truth.
- Owned end-to-end model-to-mobile pipeline — from deep learning training and validation through FastAPI deployment and mobile SDK integration — cutting release cycle time by 40%.

AI Engineer & Data Analyst | Taghna | Jeddah, Saudi Arabia

March 2025 — June 2025

- Trained end-to-end deep learning classification models on industrial sensor data achieving 88% accuracy on a 12,000-sample holdout set — covering data ingestion, feature engineering, architecture selection, and hyperparameter optimization.
- Developed automated benchmarking and real-time performance monitoring dashboards tracking model accuracy, robustness, and statistical distribution drift — enabling 3× faster anomaly detection versus manual review.
- Audited and corrected annotation quality across a 9,000-sample industrial sensor dataset with operations teams, eliminating 8.3% systematic label errors and improving production model F1 score by +5.1 points.

IoT Systems Engineer & Full-Stack Developer | Freelance / Independent Projects

2023 — 2025

- Designed and deployed multi-node IoT monitoring systems using ESP32 clusters with MQTT broker, processing 100K+ sensor events/day at sub-50ms end-to-end latency — serving 3 SME clients across Saudi Arabia.
- Engineered BLE 5.0 mesh network firmware in C++ for real-time asset tracking, achieving ±1.5-meter indoor positioning accuracy, integrated with React.js dashboards displaying live floor-map overlays.
- Developed an edge-deployed YOLOv5 plant-disease detection system achieving 94% accuracy on a 10K-image dataset, running on Raspberry Pi with sub-500ms alert latency — delivered to 2 agricultural clients in Saudi Arabia.
- Automated OTA (over-the-air) firmware updates for a fleet of 50+ ESP32 devices via AWS IoT Core + S3, reducing field-update time by 80% and eliminating manual on-site deployment entirely.
- Built React.js + Chart.js dashboards for real-time industrial sensor visualization with sub-second data refresh rates, serving 3 SME clients across Saudi Arabia.

Embedded Systems Intern | Technology Startup | Jeddah, Saudi Arabia

2024

- Wrote production STM32 HAL drivers for I2C/SPI sensor peripherals (temperature, humidity, CO2), integrated with FreeRTOS (real-time embedded OS) achieving deterministic <0.8ms interrupt response time.
- Built Python ETL pipeline ingesting 500K daily sensor readings into InfluxDB, supporting Grafana dashboards used by 6 operational staff for live facility monitoring.
- Contributed to PCB bring-up and hardware debugging using oscilloscopes, logic analyzers, and JTAG probes — reducing board bring-up time by 35% through systematic fault isolation.

KEY PROJECTS

AURA — Computer Vision Safety Platform for Industrial & Mining Environments

YOLO v8/v11 | OpenCV | FastAPI | RTSP Streaming | Python | Docker

Industrial AI safety systems prevent compliance violations and reduce injury risk at scale. AURA addresses this directly across hazardous environments.

- Architected a multi-module computer vision safety platform integrating PPE compliance detection (helmets, vests), unauthorized zone intrusion alerts, and live RTSP (IP video streaming) analysis — processing 6 simultaneous streams at 28 FPS.
- Achieved 94.2% mAP PPE detection accuracy on an 11,000-image validation set using YOLO v8/v11 fine-tuned on domain-specific industrial imagery.
- Deployed automated alerting pipelines and compliance dashboards, reducing manual safety audit overhead by 73% and enabling real-time hazard intervention.

Wathba — Predictive Mining Analytics Engine (National Competition Finalist — Future Minerals Pioneers)

Machine Learning | Predictive Analytics | Python | scikit-learn | Mining Operations

Selected as a national finalist in the Saudi government's Future Minerals Pioneers competition — validating AI talent applied to strategic mining challenges.

- Developed ML-driven predictive models for equipment failure prediction and ore-processing optimization achieving 88% prediction accuracy and 64% reduction in false alarms on a 47,000-record industry dataset.
- Advanced to the national finals from 38 competing teams, representing one of Saudi Arabia's top AI engineering solutions applied to the mining sector.

Smart Greenhouse Automation System

ESP32 | MQTT | React | FastAPI | YOLOv5 | Edge AI | Python

Reduces crop loss and water waste in arid-climate agriculture — a strategic challenge for Saudi food security.

- Deployed a 12-node ESP32 sensor mesh with automated irrigation control and computer-vision pest detection achieving 94% accuracy at sub-500ms alert latency on Raspberry Pi — processing 2,880 sensor readings/hour across 340 m² of greenhouse coverage.

BLE Indoor Asset Tracking Platform

nRF52840 | BLE 5.0 | Node.js | React | PostgreSQL | RSSI Triangulation

Delivers enterprise-grade indoor positioning at a fraction of the cost of proprietary RTLS systems.

- Built an indoor positioning system using BLE (Bluetooth Low Energy) 5.0 RSSI triangulation achieving ± 1.8 -meter accuracy, with live React dashboard and REST API supporting 200+ simultaneous device connections.

rPPG Biometric Health Engine (Production-Deployed at VLEED)

OpenCV | MediaPipe | FastAPI | Docker | Mobile SDK | Python

Contactless vital sign monitoring eliminates wearable sensors — enabling health screening at scale in clinical and consumer settings.

- Built a real-time rPPG engine extracting Heart Rate, HRV, and SpO2 from live facial video with no physical contact — achieving 91% HR accuracy (± 2.5 BPM MAE) vs. clinical pulse oximeter baseline at 28 FPS, deployed to production mobile SDK at VLEED with < 145 ms end-to-end latency.

Enterprise RAG Knowledge Pipeline

LangChain | FAISS | ChromaDB | Docker | FastAPI | OpenAI API | Vector Search

Delivers production-grade semantic search over private enterprise data without proprietary tooling.

- Productionized a Retrieval-Augmented Generation (RAG) pipeline over a 30,000-document health and enterprise knowledge base, delivering < 165 ms semantic query response latency in production.
- Implemented hybrid vector + keyword search (FAISS, ChromaDB) achieving 76% retrieval precision@5 on a domain-specific evaluation benchmark, with full FastAPI deployment and Docker containerization.

EDUCATION & CREDENTIALS

B.Sc. Artificial Intelligence — Jeddah International College, Jeddah, Saudi Arabia | Graduated 2026 | **GPA: 4.35 / 5.0**

AI Training — SDAIA | Saudi Data & Artificial Intelligence Authority | Advanced AI/ML curriculum and applied projects

AI Training — KAUST | King Abdullah University of Science & Technology | Applied AI research and engineering program

Certifications: AWS Certified Cloud Practitioner | Google IT Automation with Python (Coursera) | IoT Specialization (Coursera) | FreeRTOS Developer (Coursera)

LEADERSHIP & COMMUNITY

President — SWAICY Club | Designed and delivered Python and AI curriculum; mentored 18 members to competition placements and first engineering roles.

Head of Student Affairs — CS/IT Department | Organized 4 university hackathons and programming competitions; built a high-performance technical community of 200+ students.