MUHAMMAD HAMZA

AI Systems Engineer | Generative AI | Computer Vision | Agentic AI

Client-facing AI/ML Engineer with approximately 3 years of experience building production-grade AI systems across Agentic AI, LLM applications, Retrieval-Augmented Generation (RAG), Computer Vision, and Machine Learning, with hands-on cybersecurity engineering across VAPT, network reconnaissance, vulnerability intelligence, API/web security, and automated reporting. Combines autonomous agents, grounded retrieval, real-time vision, and security tooling to turn complex technical signals into explainable decisions, workflows, and reports. Owns delivery from client requirements and architecture through implementation, evaluation, deployment, and production support while leading a four-person development team. Published AI researcher in deep learning and medical imaging.

EMAIL
hamza.ai.official@gmail.com

WHATSAPP
+92 341 7215920

LINKEDIN
linkedin.com/in/hamzzaz

LOCATION
Lahore, Pakistan

GITHUB PORTFOLIO KAGGLE FIVERR GOOGLE SCHOLAR

EXPERIENCE

CodingCops — 2026 – Present

AI Engineer (SE)

- Build agentic AI applications using DeepAgents and LangGraph, combining task planning, tool use, and multi-step execution.
- Develop agent harnesses to coordinate execution, manage context, and connect agents with application tools and services.
- Work with Memgraph to model and query graph-based data for AI application workflows.
- Integrate Amazon DynamoDB for persistent application data and agent state management.
- Connect agent workflows with backend services and data stores to deliver end-to-end AI product functionality.
- Deploy production-grade AI projects, supporting application integration and ongoing improvements.

CyberZeus — 08/2025 – Present

CYBERZEUS SOFTWARE SYSTEMS — 01/2026 – Present

AI Engineer

- Own end-to-end delivery of production-grade AI/ML solutions, from direct client requirements and solution architecture through implementation, evaluation, deployment, and production support.
- Lead direct client discussions to understand business requirements, translate functional and technical needs into AI solutions, present demos and progress updates, gather feedback, and communicate delivery expectations.
- Lead a four-person development team of two junior developers and two interns, providing task planning, technical mentoring, code and implementation reviews, integration support, troubleshooting, and delivery coordination.

CYBERZEUS GAMES — 08/2025 – 12/2025

AI Developer

- Designed, developed, and fine-tuned AI and machine-learning models for web-based and software products, supporting model integration, evaluation, testing, and iterative improvement.
- Applied Python-based AI development and intelligent-agent concepts to product features and engineering workflows in a game/software environment.

CORE TECHNICAL SKILLS

Agentic AI & LLM Engineering
LLMs, LangChain, LangGraph, DeepAgents
Agent Harnesses, Multi-Agent Workflows, Tool Calling
State & Memory Management, MCP
Prompt Engineering

ML, Deep Learning & Computer Vision
Machine Learning, Deep Learning, NLP
YOLO, OpenCV, ByteTrack
Object Detection & Tracking
CNNs, Transformers

RAG & Retrieval Systems
RAG Architecture, Embeddings
Document Chunking & Indexing
Semantic & Vector Search
FAISS, ChromaDB, LlamaIndex
Retrieval Evaluation

Programming, Backend & Data
Python, FastAPI, REST APIs
NumPy, Pandas, SQL, PostgreSQL
Memgraph, Amazon DynamoDB

Cloud, DevOps & Deployment
AWS, Docker, Kubernetes
GitHub Actions, CI/CD, Git

Cybersecurity & VAPT
Nmap, Network Reconnaissance, Port Scanning
OS/Service Fingerprinting, DNS/DNSSEC Analysis
CVE/KEV/EPSS Enrichment
OWASP API Top 10, Web/API Assessment
SNMP, SSL/TLS Auditing, JWT/Session Security
PCAP Analysis

EDUCATION

BS Software Engineering, CGPA: 3.80
University of Okara
2020–2024

FSc Pre-Engineering
Punjab Group of Colleges

Minhaj Solutions — 04/2024 – 08/2025

AI Developer

- Built an agentic RAG-based support assistant using LLM reasoning to understand and resolve technical product queries.
- Designed the knowledge pipeline from document chunking and embeddings through vector indexing, semantic retrieval, context selection, and response generation.
- Implemented retrieval-augmented, context-aware query handling to produce real-time answers grounded in product knowledge instead of relying on fixed rule-based flows.
- Contributed to computer-vision workflows covering image preprocessing, object detection/classification, model inference, evaluation, and backend integration for client-facing AI features.

Fiverr — 07/2023 – Present

Freelance AI/ML Engineer & Python Developer

- Delivered 12+ international AI/ML and Python engagements spanning RAG applications, AI agents, workflow automation, data analysis, and computer-vision solutions for client-facing business use cases.
- Built production-oriented AI workflows using Python, FastAPI, LangChain/LangGraph, OpenAI and Gemini, vector databases, PostgreSQL, REST APIs, and Docker, with emphasis on grounded outputs and reliable integrations.
- Maintained a 4.8/5 Fiverr rating across 13 client reviews, with strong feedback for communication, delivery quality, responsiveness, and repeat collaboration.

PROJECTS

AI Agents for Cybersecurity Operations & Vulnerability Intelligence

Problem solved: fragmented network, web, API, CVE, and compliance findings force analysts to correlate evidence manually, slowing vulnerability prioritization, remediation, and reporting.

- Built a full-stack security operations platform combining LangGraph/LiteLLM agents, RAG knowledge enrichment, FastAPI orchestration, a Go/gRPC high-throughput execution engine, and a Flutter analyst desktop for scans, findings, vulnerabilities, reports, and AI-assisted analysis.
- Implemented modular security workflows for host discovery, port scanning, OS/service fingerprinting, DNS/DNSSEC analysis, CVE intelligence with KEV/SSVC/EPSS enrichment, OWASP API Top 10, XSS/web checks, SNMP, SSL/TLS, JWT/session security, PCAP forensics, and exposure/compliance assessment.
- Designed agentic tool selection and evidence correlation to combine live scan outputs with retrieved security knowledge, deduplicate and prioritize findings, generate remediation roadmaps and analyst-ready reports, and persist vulnerability history with streamed execution status.

Orbital RF Mission & Distributed Edge Operations Platform

Problem solved: fragmented satellite operations and unreliable RF execution across desktop, backend, remote agents, SDR hardware, fleet state, and mission visualization.

- Engineered the end-to-end architecture connecting a Flutter desktop client, Python APIs/services, queued commands, distributed Linux edge agents, telemetry/state, and controlled payload execution.
- Stabilized authenticated half-duplex RF payload delivery over LimeSDR, addressing identity/key integrity, RX timing, SDR process contention, false-success states, compact terminal status, Fleet lifecycle, and durable execution evidence during live M1/M2 testing.
- Integrated RF Link Analysis with 3D Space View to visualize selected targets and orbital fleet nodes, orbital context, beam/link state, contact windows, link margin, payload effects, and mission closeout.

VisionGuard AI — Real-Time Warehouse Safety Intelligence

Problem solved: conventional object detection can show boxes but does not explain whether a forklift and pedestrian are actually moving toward a dangerous near-miss.

- Built a real-time vision pipeline with YOLO detection and ByteTrack tracking, then added virtual safety zones, homography-based real-world coordinates, distance, trajectory, speed, and closing-speed analysis.
- Implemented TTC/closest-approach prediction and explainable 0–100 risk scoring using distance, motion, shared zones, exposure duration, and detection confidence.
- Generated investigation-ready evidence for critical events, including pre/post-event video clips, snapshots, track IDs, risk reasons, timelines, and structured JSON metadata.

NeuroInsight AI — Agentic Clinical Intelligence for 3D MRI (Research Prototype)

Problem solved: medical-imaging AI often stops at a prediction score, giving reviewers limited visibility into anatomical evidence, measurements, supporting research, and reasoning.

- Designed a specialized-agent workflow spanning imaging inspection, segmentation, quantitative measurements, clinical findings, RAG evidence retrieval, case reasoning, structured reporting, and human review.
- Built a Flutter product experience for multi-planar and 3D MRI inspection, quantitative biomarkers, explainability heatmaps, case-aware assistance, retrieved evidence, and review-ready reports.
- Designed case-aware RAG to combine active-study findings with research and clinical knowledge while keeping retrieved evidence distinguishable from AI-generated synthesis.

CERTIFICATIONS & ACHIEVEMENTS

Fully Funded Scholar — Design & Development of Unmanned Aerial Systems
Samara National Research University, Russia | Nomination by the University of Okara | 2025
Official University Announcement

Prime Minister's Youth Laptop Scheme — Merit Recipient
Higher Education Commission (HEC), Government of Pakistan | 2023
Achievement Post

Kaggle Notebook Expert Badge, Kaggle
Ranked Top 2,100 of ~60,000+ contributors in Kaggle Notebooks.

RESEARCH & PUBLICATIONS

Resilient UAV Network Security: Cross-Attentive Alignment of Sparse and Temporal Features for Adversarial Intrusion Detection
Vehicular Communications, 2026, Article 101094 | Available online 28/10/2026 (in press)
DOI: 10.1016/j.vehcom.2026.101094 Read research summary →
Co-authored an adversarial hybrid framework combining TabNet and a Transformer Encoder through Multi-Head Cross-Attention (MHCA), with FGSM adversarial training. Achieved 99.47% accuracy on CIC-IDS2017 and 98.24% on UNSW-NB15, improved minority Botnet recall by 33.07% over standalone baselines, and retained a 0.37–1.24 MB footprint for resource-constrained edge deployment.

3D-MobiBrainNet: Multi-class Alzheimer's disease classification using 3D brain magnetic resonance imaging
Ain Shams Engineering Journal, 16(11), 2025, Article 103714 | Published 27/08/2025
DOI: 10.1016/j.asej.2025.103714 Read research summary →
Introduced a multi-plane 3D MRI classifier using depth-wise separable feature extraction, feature recalibration, and axial/coronal/sagittal fusion; achieved 97.33% accuracy and F1-score with 99.92% AUC on the ADNI dataset while emphasizing computational efficiency.

WIDENet: A Novel Lightweight CNN for Robust Skin Lesion Segmentation
IEEE Access, Volume 14, 2026, pp. 57763–57781
DOI: 10.1109/ACCESS.2026.3683571 Read research summary →
Developed a lightweight lesion-segmentation framework combining multi-resolution wavelet decomposition, transformer-based contextual encoding, and attentive feature refinement to preserve irregular lesion boundaries; evaluated across PH2, ISIC 2017, and ISIC 2018.

Hierarchical deep learning with stability-driven normalization and regularization for accurate multi-class diabetic retinopathy diagnosis
Ain Shams Engineering Journal, 17(12), 2026, Article 104453
DOI: 10.1016/j.asej.2026.104453 Read research summary →
Developed a compact hierarchical vision encoder for five-class diabetic-retinopathy grading, addressing class imbalance, subtle retinal features, and overfitting through stability-driven normalization and regularization; achieved 98.91% accuracy on APTOS with class-wise AUCs of 0.99–1.00.