

Muhammad Mughees Ul Haq

+923361448860 | mugheesawan622@gmail.com

 [Muhammad Mughees ul haq](#)

Lahore, Punjab - 54000, Pakistan

OBJECTIVE

Machine Learning Engineer building production computer vision, deep learning, and ML pipeline systems at Wortel.ai. Skilled across the full stack: detection (YOLO, MMDetection), segmentation (SAM, YOLO11-seg), pose, depth (Depth Anything), foundation-model integration, and MLOps (ONNX, AWS, Docker). Strong record of taking models from research notebooks to evaluated, shipped pipelines.

EXPERIENCE

• Wortel.ai

June 2024 – Present

Machine Learning Engineer

Lahore, Pakistan / Remote

- Shipped a privacy face-blur pipeline combining **YOLO11m-seg** (2-class), **SAM**, and a face detector with multi-rotation gated voting; **98.3% face-bbox recall**, **91.2% F1** across a 4-pipeline evaluation framework.
- Engineered a two-stage waste-bin fill-level estimator using **Depth Anything v3 ONNX** + XGBoost regression; drove **MAE from 18.57 to 2.14** via letterbox preprocessing, regressor sweep, and temporal anchor calibration.
- Trained an **MMPose**-based 4-corner keypoint detector for static-camera bin localization; exported to **ONNX** for the depth pipeline.
- Integrated foundation models (**OWLv2**, Gemini, YOLOE embeddings) into open-vocabulary detection pipelines; benchmarked prompt strategies for cluttered-bin classification.
- Deployed and managed training / inference on **AWS EC2**; built barcode reading with **PaddleOCR** + custom **YOLOv8**.
- Collaborated with Proofminder Limited (Hungary) on agriculture automation: annotation cycles, retraining, and evaluation harnesses.

• Wortel.ai

Jan 2024 – Apr 2024

Machine Learning Intern

Lahore, Pakistan

- Trained and fine-tuned **YOLO** and **RCNN** models for object detection and classification.
- Built Python modules for data filtration, cleaning, and preprocessing across ML experiments.

EDUCATION

• Information Technology University (ITU)

Sept 2021 – May 2025

Bachelor of Science in Computer Engineering

Lahore, Pakistan

- Relevant Coursework: Data Structures & Algorithms, Artificial Intelligence, Machine Learning, Deep Learning, Computer Vision, Database Systems.

PROJECTS

• Picture-Frame Face-Blur Pipeline: Privacy segmentation for wall photos (Wortel.ai)

Tools: PyTorch, YOLO11m-seg, SAM, ONNX, Python

- Designed a 2-class frame-with-face vs without detector on ~9k images, fine-tuned with ~1.3k mined hard-negatives.
- Built a 4-pipeline benchmark on two metrics: privacy coverage and frame-polygon IoU.
- Shipped v3 frame-only as production winner: **F1 91.0%** (face-bbox), **76.4%** (frame-polygon), ~5s inference vs ~30s prior.

• Smart Car Health Inspection System: AI-powered vehicle diagnostic system (Final Year Project)

Tools: Python, YOLOv8, ESP32, Gradient Boosting, OpenCV, Flask

- Combined **YOLOv8** damage detection with OBD-II sensor analysis for unified health scoring; **84%+ mAP** and **80%+** engine health prediction, cutting inspection time by 50%.

• Smart RAG Agent: Production-grade local AI assistant

Tools: Python, FastAPI, ChromaDB, Ollama (Kimi K2.5), nomic-embed-text, Docker

- Architected a fully local **RAG agent** with dual-mode router (tool-use + JSON-classifier fallback), sub-2s response time.
- Added hardening: circuit breaker, rate limiting, structured logging; **100% pass rate** on routing benchmark.

SKILLS & CERTIFICATIONS

- Languages:** Python, C++, C, C#, SQL, JavaScript, MATLAB
- ML / DL Frameworks:** PyTorch, TensorFlow, MMDetection, MMPose, YOLOv8/11, SAM, Scikit-learn, XGBoost, Albumentations, OpenCV
- Vision & Foundation Models:** YOLO11-seg, SAM (ViT-B), OWLv2, YOLOE, Depth Anything v3, DINOv2 / CLIP, HRNet, PaddleOCR
- LLMs & RAG:** OpenAI (GPT, Vision), Gemini, Ollama (Kimi K2.5), ChromaDB, RAG, function calling, nomic-embed-text
- MLOps & Deployment:** ONNX Runtime & export, Docker, AWS (EC2, S3), FastAPI, Flask, ClearML, Git, Linux
- Certifications:** AWS Cloud Foundations, AWS Data Engineering, Fast.ai Deep Learning, CS50 Python