AMIN WAFI

AI & Machine Learning Engineer
Computer Vision | LLM Systems | Privacy

CONTACT

aminwa@icloud.com
linkedin.com/in/amin-wafi
github.com/aminwa
Athens, Greece | Open to relocation

EDUCATION

BSc (Hons) Computer Science
University of Derby
Expected 2026

TECHNICAL SKILLS

Languages
Python, C, C++, SQL, Bash

AI / ML
PyTorch, Hugging Face, YOLOv8, ONNX, OpenCV, scikit-learn, RAG, fine-tuning, LLM APIs

Web & Backend
Flask, SQLAlchemy, PostgreSQL, REST APIs, SQLite

Data
Pandas, NumPy, Matplotlib, LAPACK

Infra & Tooling
Docker, Linux, Git, CI/CD, GitHub Actions

AI Governance
GDPR, EU AI Act, EDPB, CNIL, PII handling, deployment ethics

Formal Methods
CSP-M, FDR4 model checking

LANGUAGES

English Fluent
Arabic Native
German Basic
Greek Basic

PROFILE

Final-year Computer Science student who trains and ships AI systems with a focus on how they behave in deployment. Work spans computer vision, on-device LLM tooling, and full-stack web applications. One principle runs through all of it: the user's data should not leave the user's machine. Also writes on the intersection of AI systems, security, and privacy law.

AW LABS — PERSONAL PROJECTS

shell-pilot (Published to PyPI, v1.0.0)

Python, Claude API, SQLite, Click, Rich

- Agentic terminal assistant that intercepts every shell command, blocking 14 classes of destructive operations (rm -rf, DROP TABLE, force-push, fork bombs) before execution with zero latency and no network call.
- Typo-corrects mistyped commands against every binary on $PATH using difflib, fully offline. AI explanations triggered only on failure, after secrets are redacted. 58 tests passing.

screenshield

Python, Tesseract OCR, spaCy, mss, GitHub Actions

- On-device screen monitor running at 2+ FPS, detecting 12 secret types including AWS/GCP/Azure keys, JWTs, SSNs, and Luhn-validated card numbers with Shannon-entropy false-positive filtering.
- Meeting-aware (Zoom/Teams/Meet detection), fully offline, zero telemetry. Masked local-only SQLite logging. CI green on GitHub Actions.

redact

Python, spaCy, Claude API, Typer

- Dual-engine PII redaction: spaCy NER and Claude run in parallel, spans merged and deduplicated. A compare command shows side-by-side where each engine fails, built to surface exactly where rule-based and model-based detection each break down.

ACADEMIC PROJECTS

Prohibited-Item Detection in Airport X-ray

PyTorch, YOLOv8, ONNX, Docker

- Fine-tuned YOLOv8m on OPIXray (8,885 images, five knife classes), reaching 0.924 mAP50 on a 1,776-image held-out test set. Recall target met on 4 of 5 classes.
- Ran a controlled YOLOv8s baseline on identical splits to justify model size; the +6.6 mAP50 gain was concentrated on the rarest, most-occluded class (+15 AP50). Exported to ONNX, built a resolution-agnostic inference API, containerised with Docker.

Full-Stack Web Applications (Deployed on Render)

Python, Flask, SQLAlchemy, PostgreSQL

- ElegantDine: role-based auth, reservation and order management, PostgreSQL in production, live on Render.
- Pandemic Resilience System: three role-based dashboards, encryption at rest, Werkzeug password hashing, audit logging, and RBAC.

Security Posture Evaluation: M&S Ransomware Incident

ISO 27001, NIST SP 800-61r3, GDPR

- OSINT-based reconstruction of the Scattered Spider attack chain. Mapped control failures to ISO 27001:2022 Annex A and NIST SP 800-61r3; proposed a costed 3.25M year-one remediation programme.

RESEARCH

Training LLMs on Web-Scraped Personal Data

Final-year Dissertation, University of Derby | 40 credits | 2026

- Argues GDPR is structurally incompatible with frontier-scale LLM training, synthesising empirical ML memorisation research (Carlini et al.), doctrinal EU law, and Charter fundamental-rights theory. Evaluated against EDPB Opinion 28/2024, CNIL 2025 guidance, and EU AI Act Article 53.
- Proposed a data-trusts governance model under GDPR Article 80(2) with three technical primitives: cryptographic dataset provenance, a training-data registry, and a memorisation-audit interface using membership-inference methods.
- Predicted the Court of Rome would annul the Italian Garante's 15M euro fine against OpenAI. It did, six weeks after submission.