Basel Mather

LinkedIn | GitHub | Portfolio | baselmathar@gmail.com | +962 7 9747 7786 | Amman, Jordan

Professional Summary

Full-Stack AI Engineer specializing in secure, on-premise Generative AI and Real-Time Computer Vision systems. Proven track record of deploying production-grade multi-modal pipelines (LLMs, OCR, VLMs) on bare-metal and edge infrastructure. Open-source creator with 17,000+ package downloads, combining deep expertise in MLOps, vector search, and low-latency inference optimization.

Education

- **University of Jordan** — Amman, Jordan
  Bachelor of Data Science; GPA: 3.75 ***(Ranked 1st in Data Science Class)*** — *Sep 2020 – Jun 2024*

Experience

- **Full-Stack AI Engineer** — Amman, Jordan
  Royal Hashemite Court — *Aug 2024 – Present*
  - **Architected Secure GenAI Systems:** Designed and deployed on-prem, air-gapped RAG & Agentic AI platforms and LLM architectures using vLLM for high-performance inference in confidential production environments.
  - **Engineered Autonomous AI Agents:** Developed production-grade AI agents featuring tool-calling, long-term memory, and strict validation layers using structured outputs to automate complex document intelligence tasks.
  - **Advanced Computer Vision & Spatial Analytics:** Built real-time pipelines for Object Detection, Tracking, Segmentation, and Pose Estimation; integrated via Sockets, WebSockets, and WebRTC for low-latency streaming.
  - **Cyber-Physical System Integration:** Orchestrated real-time perception and decision-making for autonomous platforms and wearable units involving LiDAR, RGB sensors, and NVIDIA Jetson/ESP32 hardware.
  - **High-Scale MLOps & Infrastructure:** Directed the setup of bare-metal AI data centers and orchestrated distributed microservices using Docker Swarm across multi-node clusters with Prometheus/Grafana monitoring.
  - **Multi-Modal Data Architecture:** Designed and maintained hybrid data systems utilizing SQL (PostgreSQL), NoSQL (MongoDB/Redis), and Vector Databases (Qdrant/Milvus) to support large-scale semantic search and retrieval.
- **Machine Learning Engineer** — Amman, Jordan
  Nafith Logistics International — *July 2023 – Aug 2024*
  - **End-to-End ML Lifecycle Management:** Spearheaded the full production pipeline for proprietary computer vision models, encompassing automated data collection, precision annotation, custom model training (PyTorch), containerization (Docker), and edge deployment.
  - **Developed Industrial ML Pipelines:** Designed and deployed end-to-end computer vision solutions for logistics, including automated object detection, tracking, and counting for high-volume cargo monitoring.
  - **Engineered Real-Time Streaming Infrastructure:** Developed high-throughput data and video pipelines using Kafka, GStreamer, and WebSockets to handle live telemetry and video feeds from remote sites.
  - **Advanced Image Processing & 3D Analytics:** Applied complex techniques including image stitching and 3D processing to enhance spatial awareness and accuracy in automated logistics environments.
  - **Scaled MLOps with Distributed Systems:** Implemented robust MLOps workflows using Docker Swarm to manage microservice-based architectures across edge and on-site deployments.
  - **Infrastructure Optimization:** Configured and tuned high-performance GPU environments using CUDA and cuDNN to accelerate deep learning training and inference cycles.

Projects

- **ClickML (AutoML Platform)**
  Designed and built a production-grade no-code AutoML platform with **5,000 users**, enabling end-to-end model training, evaluation, and deployment.

- **RetinAI (Computer Vision / Healthcare)**
  Developed RetinAI, a computer vision platform that analyzes retinal images, automates disease diagnosis, and generates reports in under a minute; awarded **2nd place** at the 8th Hakeem Academy Competition.
- **Damj (LLM Tooling – Python Package & Platform)**
  Developed an open-source Python package (**9K+ downloads**) for structured prompt engineering and project context aggregation; extended into a containerized, no-code platform deployed via Docker.
- **EmotionLens (Real-Time Computer Vision)**
  Built a real-time deep learning system for face detection, tracking, and emotion recognition using YOLO-based object detection and multi-stage vision pipelines.
- **PHAI (GenAI Healthcare Platform)**
  Implemented an LLM-powered platform using RAG to help users understand medical prescriptions, including document ingestion, semantic search, and conversational interfaces for accurate information retrieval.
- **Streamlit Arabic Support Wrapper (Python Package)**
  Created an open-source Python package for Streamlit component with **8K+ downloads** enabling correct Arabic text alignment and rendering.
- **Search Engine for Arabic Content (NLP)**
  Built a full-stack Arabic search engine from 500 articles, including text preprocessing, indexing, ranking, and a web-based user interface.

Core Skills

Generative AI & ML — LLMs (local: vLLM, Hugging Face, Ollama; cloud: OpenAI, Gemini), RAG & AI Agents, Vision-Language Models (VLMs), PyTorch, TensorFlow, Scikit-learn.

Computer Vision — Image Recognition, Object Detection, Multi-Object Tracking, Pose Estimation, Segmentation, OCR, Real-Time Video Analytics, Edge Deployment.

Data Architecture — Vector Databases (Qdrant, Milvus), SQL (PostgreSQL, MySQL), NoSQL (MongoDB, Redis), Object Storage (Garage/S3), Data Streaming (Kafka, GStreamer).

Infrastructure & MLOps — Docker, Docker Compose, Docker Swarm, Nginx, Traefik, CI/CD Pipelines, Bare-metal Deployment, On-premise Infrastructure, Private Docker Registry.

Backend & Protocols — Python, FastAPI, REST APIs, WebRTC, Sockets, WebSockets.

Monitoring & Tools — Linux, Git, Prometheus, Grafana, Cron Jobs.

Certifications

- Docker Certified Associate (LearnKartS)
- Deep Learning Specialization (DeepLearning.AI)
- NLP Specialization (DeepLearning.AI)

Volunteering & Leadership

- Director of the ACM Magazine at JU Chapter and Contributing Writer
- Technical Team Lead at IEEE Computational Intelligence Society (CIS)

Languages

- Arabic (Native)
- English (Professional)