

# GONÇALO GUEDES

Porto, Portugal (Open to Global Relocation) • +351 912161011 •  
goncalo.manuelguedes@gmail.com • [LinkedIn](#) • [goncalo-guedes.com](http://goncalo-guedes.com)

---

## PROFESSIONAL SUMMARY

Senior Machine Learning Engineer and Published Researcher with 4+ years of experience specializing in **high-performance inference optimization** and **Edge AI**. Expert in architecting scalable systems (15M+ monthly transactions) using **NVIDIA Triton, TensorRT, and ONNX**. Proven track record of leading engineering teams and transitioning complex R&D projects into production-grade hardware and software products.

---

## TECHNICAL SKILLS

- **Deep Learning & Inference:** PyTorch, TensorFlow, NVIDIA TensorRT, Triton Inference Server, ONNX, TensorFlow Lite, CoreML, Model Quantization (INT8/INT4).
  - **GenAI & NLP:** RAG (LangChain, OpenAI, Weaviate), Transformers, Entity Extraction, Vector Databases, Multi Agents System.
  - **Computer Vision & Imaging:** Image Classification, Segmentation, OCR, Edge AI, X-ray tomography, Hyperspectral Imaging.
  - **Cloud & Infrastructure:** **AWS (Certified)**, Azure, Docker, Kubernetes, CI/CD, Git, SQL.
- 

## PROFESSIONAL EXPERIENCE

**Glintt Global | Lead of AI/ML Engineering Chapter | Dec 2025 – Present**

- **Direct a cross-functional chapter of 9 MLEs**, managing the technical roadmap and delivery for a diverse portfolio including RAG-based agents, Computer Vision, and NLP solutions.
- **Orchestrate project lifecycles** from initial R&D to production-ready solutions, ensuring technical excellence and alignment with enterprise client KPIs.
- **Mentor and scale the engineering team**, establishing best practices for code quality and performance optimization to ensure high-throughput inference standards (Triton/TensorRT).
- **Bridge the gap between business and tech**, translating complex AI requirements into scalable architectural designs for large-scale industrial and mail-service clients.

## Glintt Global | Senior Machine Learning Engineer | Apr 2025 – Dec 2025

- **Designed a multi-modal CV/NLP system** for a national mail carrier to automate entity extraction from **15M+ letters per month**.
- **Boosted extraction accuracy from 76.9% to 98.7%** by replacing a legacy system with a custom Deep Learning pipeline for real-time address validation.
- **Optimized inference throughput** using **NVIDIA Triton and TensorRT**, slashing compute overhead and ensuring scalability for peak seasonal traffic.

## Glintt Global (Glintt Next) | Machine Learning Engineer | Jun 2024 – Apr 2025

- **Developed a production-grade RAG agent** using **OpenAI, LangChain, and Weaviate**, optimizing document indexing and grounding for high-precision information retrieval.

## Cork Supply | Machine Learning Engineer | Feb 2022 – Jun 2024

- **Pioneered the "X100"**, an industry-first AI-driven X-ray tomography system used to classify cork quality and oxygen transfer rates non-invasively.
- **Eliminated 3-month destructive lab testing cycles** by providing real-time quality classification through an automated Deep Learning system.
- **Engineered for the Edge** using **TensorFlow Lite**, achieving a real-time throughput of **6 inferences per second** on resource-constrained hardware.

---

## PUBLICATIONS

- **Evaluating cork coating homogeneity using hyperspectral imaging and texture analysis.** *Journal of Near Infrared Spectroscopy*, 2025.
- **Handheld Near-Infrared Spectroscopy: State-of-the-Art Instrumentation and Applications.** *Chemosensors*, 2023.

---

## EDUCATION & CERTIFICATIONS

- **University College Dublin** | Research Scholar (SensorFint Grant/COST Action), 2024.
- **Universidade Católica Portuguesa** | **M.Sc. in Biomedical Engineering**, 2023.
- **Universidade Católica Portuguesa** | **B.Sc. in Bioengineering**, 2020.
- **AWS Certified Cloud Practitioner** (2025 – 2028).
- **PCEP™** – Certified Entry-Level Python Programmer.