

Daniel Amirifard

AI/ML Engineer

Milan, Italy | +39 379 208 2934 | daniel.amirifard@gmail.com | linkedin.com/in/daniel-amirifard | github.com/daniel-amirifard | daniel-amirifard.github.io

Professional Summary

AI/ML Engineer with 4+ years of experience developing Machine Learning, Deep Learning, and Generative AI solutions. Graduated from **Politecnico di Milano** and contributed to **30+ AI/ML projects** across academic and industry settings, including **RAG systems, LLM applications, NLP, computer vision, time-series forecasting, agentic workflows, and model deployment**. Skilled in building scalable, production-ready AI systems, with project experience involving **STMicroelectronics, Nokia, and Baker Hughes**.

Skills

Programming: Python, SQL

Machine Learning & AI: TensorFlow, PyTorch, Keras, Scikit-learn, Pandas, NumPy, SciPy, Matplotlib

Generative AI & LLMs: RAG, Agentic AI, LangChain, Prompt Engineering, Chatbots

Cloud & DevOps: AWS (Amazon Bedrock, S3), Azure ML Studio, Linux, Git, GitHub, Docker, FastAPI, CI/CD

Others: English (C1), Italian (A2), Teamwork, Organization, Communication, Creativity

Work Experience

Artificial Intelligence Engineer

11/2025 – Present

Hermes/Baker Hughes, Full-time

Milan, Italy

- Designed **3+ production-oriented** GenAI solutions from PoC to deployment, using RAG pipelines, LangGraph, and Amazon Bedrock to improve automation, scalability, and business impact.
- Reduced manual invoice review effort by **~85%** by automating milestone classification from structured and unstructured invoice text.
- Delivered AI applications across **3 industry domains**, including financial invoice classification, chemical-sector workflows, and energy-sector applications, with one ML classification pipeline achieving a 93% F1-score.

Machine Learning Engineer

05/2025 – 11/2025

STMicroelectronics, Full-time

Milan, Italy

- Designed and deployed automated **neural architecture search (NAS)** workflows for **LLM agent bots**, improving model accuracy and computational efficiency by 40%.
- Implemented zero-cost proxy methods for **training-free** architecture evaluation, reducing GenAI benchmarking time by **10x**.
- Constructed and integrated an **AutoNN LLM agent**, that automatically designs neural network structures, reducing model architecture generation time 70%.

Machine Learning Engineer

02/2023 – 10/2024

Polimi Data Scientists Association, Part-time

Milan, Italy

- Developed, trained, and **optimized 10+** machine learning models in Computer vision, anomaly detection and time series.
- Organized data **pipelines** for efficient **ETL** processes and sourced relevant datasets, used across 5+ projects including SIIT and Breast Cancer detection.
- Presented technical progress in biweekly sessions and contributed to the completion of **4 major student-led AI** projects.

Data Scientist

09/2021 - 09/2022

Doran Smart-City, Full-time

Tehran, Iran

- Developed Random Forest and XGBoost models to analyze urban traffic patterns, supporting traffic-flow optimization with a reported **15% improvement**.
- Redesigned predictive models for smart energy management, contributing to a **10% reduction** in infrastructure energy consumption.

Education

Politecnico di Milano M.Sc. in Telecommunication Engineering – Computer Science and Communications	09/2022 - 03/2026 Milan, Italy
Amirkabir University of Technology B.Sc. Engineering	09/2016 - 09/2021 Tehran, Iran

Projects

- RAG-Based Financial Invoice Classification Pipeline** | Amazon Bedrock, GenAI, Agentic AI systems, LangGraph
- Built an AI-powered invoice classification pipeline using RAG, FAISS, LangGraph, and Mistral via Amazon Bedrock, automating milestone classification from invoice text fields and achieving a 93% F1-score.
- Long Form Question Answering ChatBot** | Retrieval-augmented generation (RAG), LLAMA, LoRA, BART, NLP, Hugging Face
- Created an NLP pipeline for Long Form Question Answering using transformers (BART, LLAMA) and RAG in a team of 4, securing a ROUGE score of 22.
- Zero-Cost Proxies in AutoNN LLM Agent** | Google Cloud, Azure ML Studio, GenAI, Fine Tuning, Agentic AI Systems, OpenAI
- Integrated zero-cost proxies into AutoNN LLM agent, achieving 0.78 Kendall’s Tau on CIFAR-x datasets for reliable training-free architecture performance estimation.
- SIIT ChatBot** | RAG, GPT, BERT, LLAMA, LangChain, PyTorch, FAISS, Prompt Engineering
- Spearheaded a technical team of 10 in the SIIT project with JEMP startup for SIIT Company, improving pharmaceutical synthesis and slashing production time by ~25% by developing a chatbot.
- Cooperative Perception In Automated Cars** | 3D Object detection, GenAI, Transformers
- Constructed a V2V cooperative perception pipeline using OPV2V, LiDAR, PointPillars, ROI filtering, and attention fusion, achieving 0.914 AP@0.5 and 0.837 AP@0.7 for 3D object detection.

Workshop Experience

Impact of non-Gaussian noise on LDPC performance Project Nokia Company	06/2024 Vimercate, Italy
Presented to over 100 attendees at Nokia, showcasing a simulation of non-Gaussian noise effects on LDPC performance using QAM 4, 16, and 64.	
MultiNet-Route-Optimizer Project Nokia Company	03/2024 Vimercate, Italy
Presented a 65% improvement in routing efficiency for a Nokia-defined problem, achieved by a 4-person team using Top-Down and Bottom-Up optimization strategies.	

Certifications

- Neural Networks and Deep Learning** | DeepLearning.AI
- Supervised Machine Learning: Regression and Classification** | DeepLearning.AI
- Linux LPIC-1** | Tehran Institute of Technology
- Python Programming** | Tehran Institute of Technology