Michele **Bellitti**

Profile

Machine Learning Engineer with hands-on experience designing and shipping end-to-end ML and GenAI solutions, from data ingestion to cloud deployment.

Comfortable across Python, SQL and modern data and cloud platforms, with prior experience in high-performance data engineering using Rust and Polars. Enthusiastic about production-grade AI systems and the boundary between classical ML, agentic GenAI and large-scale data processing.

Contact details

@ michelebellitti272@gmail.com
Github Page
Portfolio Website
LinkedIn

Personal information

Citizenship: **Italy**

Family: **Not Married**

Languages: **English** (C1), **Italian** (native)

Skills

- **Languages & ML:** Python, SQL, Polars, PyTorch, NetworkX
- **GenAI:** LangChain, LangGraph, Langfuse, MCP, RAG, tool-calling, OpenAI / Azure OpenAI
- **Azure:** Blob Storage, Functions, Cosmos DB, Storage Queue, AI Foundry, Document Intelligence
- **GCP:** Cloud Run, Cloud Scheduler, IAM / service accounts
- **Data & MLOps:** Snowflake, DBT, PySpark, Pydantic, Docker, Git, CI/CD
- **Other:** Rust, C/C++, FastAPI, Streamlit, Linux
- Cross-functional collaboration, technical writing

Experience

machine learning engineer at Prometeia **May 2025 – Present**

- Designed and contributed to GenAI solutions (RAG, agentic workflows with LangGraph, MCP-based tooling) integrating LLMs into existing credit-risk pipelines.
- Designed and shipped a scalable monitoring solution for Sustainable AI.
- Designed and implemented end-to-end ETL pipelines for transactional data using Azure, PySpark and Python.
- Designed and shipped a Spark-driven data quality and transformation library for data ingestion and processing at scale.
- Productionised ML and GenAI components with microservice architectures, containerised deployments and cloud-native observability.

Data Engineering Intern at Axyon AI **September 2024 – April 2025**

- Profiled the in-house dataset-generation software, identifying bottlenecks and quantifying performance limits in detailed reports.
- Contributed to the next-generation rewrite using Polars and Rust, materially improving throughput and scalability on large datasets.
- Built a DBT-based ingestion project consolidating multiple data providers into a centralised Snowflake data warehouse.

Education

M.Sc. in Computer Science Engineering Large Scale Artificial Intelligence Engineering *University of Modena and Reggio Emilia*. **2022–2026 (expected)**

Bachelor of Science in Computer Science Engineering. Department of Engineering Enzo Ferrari. *University of Modena and Reggio Emilia*. **2019–2022**

◇ Thesis title: *Enhancing traditional Congestion Control through Machine Learning Techniques*

Additional education

Machine Learning in Production *Coursera.* **2024**

◇ End-to-end production ML system design: project definition, data needs, modeling strategies, templates and deployment technologies.

Full-Stack Web Development: HTML, CSS, Javascript, Node, React, PostgreSQL *Udemy* **2025**

◇ Comprehensive full-stack web development course covering modern tools and technologies.