

MATTIA CURRI

Master's Degree Computer Science Student - Software Developer - AI Engineer

✉ mattia2002@gmail.com

🔗 [mattiacurri](https://github.com/mattiacurri)

🌐 [mattiacurri](https://mattiacurri.github.io)

🌐 mattiacurri.github.io

ABOUT ME

Ciao! I'm Mattia, a Computer Science Master's student at the University of Bari. My work spans VLM distillation, hybrid RAG architectures, and, since plain NLP wasn't enough, teaching language models to solve funny games like crosswords. I thrive at the intersection of deep learning, natural language processing and computer vision, always curious about the next hard problem.

EDUCATION

Master's Degree in Computer Science - Artificial Intelligence Track

University of Bari, Italy

📅 110/110 (with honours and special mention for academic career)

📅 2024 - 2026

- Master's thesis at the [SWAP](#) research group: *A Hybrid Retrieval-Augmented Generation Pipeline Combining Structured Records and Knowledge Graphs for Public Procurement*
 - Combines SQL retrieval over 103M+ ANAC records with graph-based retrieval over procurement documents; evaluated on 40 questions across 18 tenders.
- Relevant courses taken: Numerical Methods, [Database II](#), [Formal Methods](#), Natural Language Processing, Software Engineering for AI-Enabled Systems, Deep Learning, Computer Vision, Quantum Computing, Machine Learning, Big Data

Bachelor's Degree in Computer Science

University of Bari, Italy

📅 110/110 (with honours)

📅 2021 - 2024

- Thesis in Computational Intelligence: *Supervised Learning Techniques for Semantic Segmentation of Aerial Images of Agricultural Fields*
 - 3 Months research internship at the [Cilab](#) Lab of the University of Bari under the supervision of Prof. [Gennaro Vessio](#).
 - Comparison of Convolutional and Transformer architectures to tackle the Supervised Task of the Agriculture-Vision Challenge at [CVPR 2024](#), training models on ~ 50k images.
- Relevant courses taken: Computational Intelligence, Algorithms and Data Structures, Databases, Software Engineering, Knowledge Engineering, Information Retrieval, Human-Computer Interaction, Computer Networks.

PUBLICATIONS

A. Porcelli, F. Di Gravina, E. Fontana, [M. Curri](#), F. D. Di Gregorio. "FFT-UniBa at Cruciverb-IT: Special Length Tokens and CSP for Italian Crossword Solving." In [Proceedings of the 9th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. EVALITA 2026](#), <https://ceur-ws.org/Vol-4195/45.pdf>

- Fine-tuning of Italian T5 with length-constrained generation and CSP-based grid solver; MRR 0.63, 34% full grid accuracy.

M. Borgo, [M. Curri](#), O. Zaroual, P. Serio, G. Pisaneschi, M. G. C. A. Cimino. "Weather-Aware Multimodal Sensor Fusion for Robust Perception." In [Proceedings of the 2nd IEEE International Conference on Application of Information Technologies in Engineering, Management and Science \(ICAI-TEMS 2026\)](#), Pisa, Italy, 10-13 November 2026.

- Camera-radar fusion for vehicle detection under adverse weather on the RADIATE dataset.
- Research conducted as part of the UIS Research Centre's [Students for Italy's Future](#) project.

WORK EXPERIENCE

Didactic Tutor (150 Hours Program)

🔗 C Teaching Formal Languages

📅 2025-2026

- Delivered tutoring sessions for first-year Computer Science students in Programming and Formal Languages courses.
- Supported students with problem-solving and exam preparation through structured meetings across the semester.

PROJECTS

Tender Copilot: Hybrid Graph-RAG for Public Procurement

Master's Thesis Project in AI & NLP

🔗 Python Graph RAG HippoRAG Neo4j SQLite FastAPI MLflow Docker LM Studio

🔗 [TenderCopilot](#)

- Architected an ingestion engine combining national **ANAC Open Data (103M+ records)** across 23 relational tables) with administrative tender PDFs parsed via **Docling OCR**, table normalization, and section-aware chunking
- Constructed an open-schema knowledge graph (**AutoSchemaKG**) via a 3-pass LLM pipeline (entity-entity, event-entity, temporal/causal event edges) with dense synonym bridging and full chunk provenance using open-weight models
- Implemented a hybrid retrieval pipeline coupling **deterministic relational scoping** (SQLite FTS5 + Text-to-SQL) with a **HippoRAG2-style graph retriever** using **Personalized PageRank (PPR)** and cross-encoder reranking

- Engineered a **100% local-first architecture** with zero cloud API reliance, deploying open-weight LLMs (Qwen, Gemma, Muse) via an API (through LM Studio) and deploying **Open WebUI** in **Docker**

Explainability-Guided Distillation for Vision-Language Models - Solo Project

Computer Vision Exam Project

</> Python PyTorch Hugging Face Grad-CAM Qwen-VL

 [ExplanationGuidedDistillation](#)

- Developed an end-to-end pipeline to distill the visual module of **Qwen/Qwen2.5-VL-3B-Instruct** into a **ViT** student
- Implemented an **explainability-guided distillation loss** based on **Grad-CAM**
- Built an automated **teacher-vs-student evaluation** with generated responses, **BERTScore**, **LLM/VLM-as-a-judge**, and blind A/B comparison on 500 images

End-to-End Code Comment Classification System - Team Project

Software Engineering for AI Exam Project

</> Python CatBoost Grafana FastAPI Docker MLflow Prometheus Github Actions

 [CodeCommentClassification](#)

- Architected a production-grade multi-label classification system categorizing code comments across Java, Python, and Pharo
- Engineered end-to-end ML pipeline: data versioning (DVC) → automated preprocessing → dual-model training (CatBoost + SetFit) → experiment tracking (MLflow) → containerized deployment
- Implemented comprehensive **data quality framework** improving dataset integrity by 12%
- Deployed full-stack solution: **FastAPI** backend with async processing → **Gradio** frontend → **Prometheus/Grafana** monitoring, with sub-2s inference latency
- Achieved **92.7% test coverage** across unit, integration, and behavioral test suites, ensuring production reliability

Corporate Credit Rating Prediction - Solo Project

AI Project

</> Python Scikit-learn Pandas Pgmypy Imblearn Matplotlib Numpy

 [Rating Prediction](#)

- Dealed with an unbalanced dataset using **Class Weights**, **SMOTE**, **ADASYN**, **SMOTETomek**, **SMOTEENN**
- Implemented a **Bayesian Network** to do probabilistic queries, with a correlation score based on *Balanced Accuracy* of **0.59**
- Achieved up to **86%** of **Balanced Accuracy** using XGBoost and Random Forest

AWARDS, HACKATHONS & COMPETITIONS

Leonardo Hackathon – Space Edition

</> AI Automation Robotics Space

 2026

- Space hackathon organized by Leonardo in collaboration with Talent Garden, with focus on AI and automations for space applications.
- Achieved 3rd place** in team.
- Selected as one of the 50 finalists** among 400 candidates
- Obtained 2000€ prize for the team.

First Italian University AI Competition

</> AI LLM Open Data

 2026

- Achieved 2nd place** in the first stage of the First Italian University AI Competition organized by [AI2B](#), out of 20 teams.
- The competition involved using open data to create a solution for the city of Bari, specifically focused on the problem of logistics and intelligent traffic management.

ITADData Hack 2025

</> Python CatBoost Scikit-Learn Imblearn Pandas

 2025

- Anomaly Detection task in Hadoop Distributed File System Logs.
- Secured 2nd place** in the **ITADData Hack 2025 competition** out of 15 teams.

AI2B Hackathon Winner

</> Cryptography Prompt Engineering

 2025

- Winning team of the AI2B Hackathon organized by [University of Bari](#) and [AI2B](#), focused on AI and cybersecurity.
- Obtained €500 winning first place out of 10 teams.

CyberChallenge 2024

</> Web Binary Exploitation Cryptography Network

 2024

- Qualified to the **national stage** by ranking in the top 6 of the local venue, reaching 21st place out of 43 teams at the national competition.
- Qualified for the local stage by being in the top 20 of the selection.

ADDITIONAL

Languages: Italian (native), English (B2 Cambridge)