

# Razvan Florian Vasile

Cesena, Italy · [florian.rvasile@gmail.com](mailto:florian.rvasile@gmail.com) · [Github](#) · [Website](#) · [LinkedIn](#)

## EDUCATION

---

### **Master's Degree in Computer Science** 2022 - 2025

University of Bologna (Unibo), Bologna, Italy — Awarded March 2026

- Completed a Master's degree in Computer Science with a focus on Machine Learning.
- Final grade: 110/110 with honours.

### **BSc (Hons) Computer Science and Mathematics** 2014 - 2018

University of Manchester (UoM), Manchester, UK — Incl. Industrial Experience Year

- Finished a joint honours degree in Computer Science and Mathematics.
- Final grade: First Class Honours.
- Awarded the Manchester Scholarship on entry (A\*AA equivalent).

## EMPLOYMENT AND EXPERIENCE

---

### **Software Engineer** 2018 - 2020 / 2020 - 2022

Onit Group / Agronica, Cesena, Italy — European Digital Agriculture 4.0

Python · SQL · NoSQL · IoT · Full-stack developer

- Developed Raspberry Pi IoT integrations in Python for warehouse equipment control.
- Full-stack development across application services supporting SQL/NoSQL databases.

### **Industrial Placement Year** 2016 - 2017

Diamond Light Source, UK — National Synchrotron facility

Python · MATLAB · Scientific software · Accelerator controls

- Developed a unified library for accelerator controls, creating a single Python interface to access previously fragmented MATLAB-based accelerator components.
- Enabled programmatic experimentation in Python, allowing precise control of magnet steering positions to optimize beam trajectories for diverse experimental setups.

### **Peer Assisted Study Sessions (PASS) Leader** 2015 - 2016

University of Manchester Teaching Assistant, UK

- Assisted first-year students with programming concepts and coursework problems.
- Organised review sessions and prepared practice materials for exams.

### **Hornet Representative** 2014 - 2016 / 2017 - 2018

University of Manchester IT Services, UK

- Resolved wired and wireless connectivity issues across university halls of residence.
- Visited residents on site to diagnose and resolve network connectivity problems.

## TECHNICAL SKILLS

---

**Languages:** Python, C++, MATLAB, JavaScript/TypeScript.

**Experienced with:** Distillation, SFT, RLVR/GSPO, Tool-Using Agents, RAG, Embeddings & Vector Search, LoRA/PEFT, Multimodal Learning, Distributed Training, GPU Programming.

**Math. Background:** Linear Algebra, Matrix Analysis, Convex Optimization, Discrete Mathematics, Propositional & Predicate Logic (Number Theory), Logic and Modelling, [\[full catalog\]](#).

References and more information available upon request.

## PROJECTS

---

### **Master's Thesis: Distilling Neurosymbolic Reasoning for Linear Algebra in Small Language Models** 2025 – 2026

[GitHub](#) · [Report](#) · [Poster](#) · [Demo](#) · [HF Space](#)

- Developed an end-to-end neurosymbolic pipeline for training a small pre-trained language model to solve linear algebra problems via sequential tool calls.
- Engineered a synthetic data generator and knowledge distillation pipeline, employing Supervised Fine-Tuning (SFT) and policy optimization algorithms such as GSPO.
- Achieved 90.26% end-to-end accuracy on the test set using \$75 worth of cloud compute.

### **Bachelor's Thesis: Solving Sudoku using Propositional Logic** 2017 – 2018

[GitHub](#) · [Report](#) · [Presentation](#)

- Implemented a CDCL SAT solver in C++ including techniques for backjumping, conflict-driven clause learning, two encoding approaches and two separate decision heuristics.
- Developed an interactive visualizer to inspect the resolution process, displaying implication graphs, unit propagations, conflicts, and backtracking steps in real-time.
- Evaluated solver performance on difficult SAT instances via VSIDS/LIS decision heuristic analysis, flame-graph profiling, and minimal versus extended encoding comparisons.

### **Distributed and Parallel Techniques for Deep Neural Networks** 2025

[GitHub](#) · [Report](#) · [Poster](#)

- Implemented GPU training experiments with cuDNN and cuBLAS, covering forward and backward passes, convolution primitives, manual kernels, and linear layers.
- Built reproducible distributed-training examples with PyTorch DDP and Docker to simulate multi-GPU and multi-node workflows on a single NVIDIA GPU.
- Conducted a literature review of distributed training and GPU libraries, comparing parallelization strategies, communication patterns, and implementation trade-offs.

### **Visual QA: Using Generative Models on Classification Tasks** 2024

[GitHub](#) · [Report](#) · [Poster](#) · [W&B](#)

- Adapted the BLIP-2 multimodal architecture for discriminative tasks by modifying the Q-Former component with a classification head.
- Fine-tuned the model with LoRA, achieving 91% accuracy on Easy-VQA and 78% on DAQUAR datasets, outperforming generative baselines.

### **Git Inspector: Querying GitHub Repositories with Local LLMs** 2025

[GitHub](#) · [Report](#) · [Poster](#) · [Demo](#) · [Website](#)

- Built Git Inspector, a Scala-based RAG system for querying GitHub repositories using semantic code retrieval with Qdrant and local open-source LLMs via Ollama.
- Implemented repository cloning, indexing, semantic chunking, and a REST API with Gradio and Scala.js front-end interfaces for code and documentation queries.
- Validated the system through acceptance tests, traceability documentation, and a usability study that achieved an average SUS score of 85% across five users.

### **Librarian Assistant** 2024

[GitHub](#) · [Report](#) · [Poster](#) · [Presentation](#)

- Built Librarian, a local question-answering assistant for personal PDF libraries using open-source LLMs and retrieval-augmented generation over book contents.
- Implemented a structured reasoning agent with vector retrieval and citation-grounded answers, supporting both local and serverless inference via Ollama and HuggingFace.

Expanded project information [available here](#).